

6-11-2020

Universal Algebraic Methods for Constraint Satisfaction Problems

Clifford Bergman

Iowa State University, cbergman@iastate.edu

William Demeo

Charles University in Prague

Follow this and additional works at: https://lib.dr.iastate.edu/math_pubs



Part of the [Algebra Commons](#)

The complete bibliographic information for this item can be found at https://lib.dr.iastate.edu/math_pubs/191. For information on how to cite this item, please visit <http://lib.dr.iastate.edu/howtocite.html>.

This Article is brought to you for free and open access by the Mathematics at Iowa State University Digital Repository. It has been accepted for inclusion in Mathematics Publications by an authorized administrator of Iowa State University Digital Repository. For more information, please contact digirep@iastate.edu.

Universal Algebraic Methods for Constraint Satisfaction Problems

Abstract

After substantial progress over the last 15 years, the "algebraic CSP-dichotomy conjecture" reduces to the following: every local constraint satisfaction problem (CSP) associated with a finite idempotent algebra is tractable if and only if the algebra has a Taylor term operation. Despite the tremendous achievements in this area (including recently announce proofs of the general conjecture), there remain examples of small algebras with just a single binary operation whose CSP resists direct classification as either tractable or NP-complete using known methods. In this paper we present some new methods for approaching such problems, with particular focus on those techniques that help us attack the class of finite algebras known as "commutative idempotent binars" (CIBs). We demonstrate the utility of these methods by using them to prove that every CIB of cardinality at most 4 yields a tractable CSP.

Keywords

constraint satisfaction problem, universal algebra, dichotomy conjecture, Mal'tsev condition, Taylor term, absorbing subalgebra

Disciplines

Algebra

Comments

This is a pre-print of the article Bergman, Clifford, and William DeMeo. "Universal Algebraic Methods for Constraint Satisfaction Problems." *arXiv preprint arXiv:1611.02867* (2020). Posted with permission.

Creative Commons License



This work is licensed under a [Creative Commons Attribution 4.0 International License](https://creativecommons.org/licenses/by/4.0/).

**UNIVERSAL ALGEBRAIC METHODS FOR
CONSTRAINT SATISFACTION PROBLEMS
WITH APPLICATIONS TO COMMUTATIVE IDEMPOTENT BINARS**

CLIFFORD BERGMAN AND WILLIAM DEMEO

Department of Mathematics, Iowa State University in Ames, Iowa, USA
e-mail address: cbergman@iastate.edu

Department of Algebra, Charles University in Prague, Czech Republic
e-mail address: williamdemeo@gmail.com

ABSTRACT. After substantial progress over the last 15 years, the “algebraic CSP-dichotomy conjecture” reduces to the following: every local constraint satisfaction problem (CSP) associated with a finite idempotent algebra is tractable if and only if the algebra has a Taylor term operation. Despite the tremendous achievements in this area (including recently announce proofs of the general conjecture), there remain examples of small algebras with just a single binary operation whose CSP resists direct classification as either tractable or NP-complete using known methods. In this paper we present some new methods for approaching such problems, with particular focus on those techniques that help us attack the class of finite algebras known as “commutative idempotent binars” (CIBs). We demonstrate the utility of these methods by using them to prove that every CIB of cardinality at most 4 yields a tractable CSP.

1. INTRODUCTION

The “CSP-dichotomy conjecture” of Tomás Feder and Moshe Vardi [FV99] asserts that every constraint satisfaction problem (CSP) over a fixed finite constraint language is either NP-complete or tractable. (The current status of the conjecture is discussed in the note at the end of this introduction.)

A discovery of Jeavons, Cohen and Gyssens in [JCG97]—later refined by Bulatov, Jeavons and Krokhin in [BJK05]—was the ability to transfer the question of the complexity of the CSP over a set of relations to a question of algebra. Specifically, these authors showed that the complexity of any particular CSP depends solely on the *polymorphisms* of the constraint relations, that is, the functions preserving all the constraints. The transfer to universal algebra was made complete by Bulatov, Jeavons, and Krokhin in recognizing that

2012 ACM CCS: [Theory of computation]: Algebraic complexity theory, Complexity theory and logic.

2010 Mathematics Subject Classification: 08A70; 03C05, 08A30, 08A40.

Key words and phrases: constraint satisfaction problem, universal algebra, dichotomy conjecture, Mal'tsev condition, Taylor term, absorbing subalgebra.

Research of both authors partially supported by National Science Foundation grant no. 1500218.

with any set \mathcal{R} of constraint relations one can associate an algebra $\mathbf{A}(\mathcal{R})$ whose operations consist of the polymorphisms of \mathcal{R} .

Following this, the CSP-dichotomy conjecture of Feder and Vardi was recast as a universal algebra problem once it was recognized that the conjectured sharp dividing line between those CSPs that are NP-complete and those that are tractable was seen to depend upon universal algebraic properties of the associated algebra. One such property is the existence of a Taylor term (defined in Section 2.2.2). Roughly speaking, the “algebraic CSP-dichotomy conjecture” is the following: The CSP associated with a finite idempotent algebra is tractable if and only if the algebra has a Taylor term operation in its clone. We state this more precisely (postponing technical definitions) as follows:

Conjecture 1.1. If \mathbf{A} is an algebra with a Taylor term in its clone and if \mathcal{R} is a finite set of relations compatible with \mathbf{A} , then $\text{CSP}(\mathcal{R})$ is tractable. Conversely, if \mathbf{A} is an idempotent algebra with no Taylor term in its clone, then there exists a finite set \mathcal{R} of relations such that $\text{CSP}(\mathcal{R})$ is NP-complete.

The second sentence of the conjecture was already established in [BJK05] and, as mentioned below, the proof of the statement in the first sentence was recently announced independently by Andrei Bulatov [Bul17] and Dmitriy Zhuk [Zhu17]. The purpose of the present work is to provide new methods for establishing the truth of first sentence of the conjecture for certain special classes of algebras. Another goal is to give concrete examples demonstrating how these methods work.

Algebraists have identified two quite different techniques for proving that a finite idempotent algebra is tractable. One, often called the “local consistency algorithm,” works for any finite algebra lying in an idempotent, congruence-meet-semidistributive variety (SD- \wedge for short). See [BK09] or [BK12] for details. The other, informally called the “few subpowers technique,” applies to any finite idempotent algebra possessing an edge term [IMM⁺10]. Definitions of these terms appear in Section 6.3. While these two algorithms cover a wide class of interesting algebras, they were not enough to resolve the full dichotomy conjecture above. Several researchers attempted to combine the two existing approaches in a way that captures the outstanding cases. In this paper, we do as well.

For example, suppose that \mathbf{A} is a finite, idempotent algebra possessing a congruence, θ , such that \mathbf{A}/θ lies in an SD- \wedge variety and every congruence class of θ has an edge term. It is not hard to show that \mathbf{A} has a Taylor term. Can local consistency be combined with few subpowers to prove that \mathbf{A} is tractable?

We can formalize this idea as follows. Let \mathcal{V} and \mathcal{W} be idempotent varieties. The *Mal’tsev product* of \mathcal{V} and \mathcal{W} is the class

$$\mathcal{V} \circ \mathcal{W} = \{ \mathbf{A} : (\exists \theta \in \text{Con}(\mathbf{A})) \ \mathbf{A}/\theta \in \mathcal{W} \ \& \ (\forall a \in A) \ a/\theta \in \mathcal{V} \}.$$

The Mal’tsev product $\mathcal{V} \circ \mathcal{W}$ is always an idempotent quasivariety, but is generally not closed under homomorphic images. A worthwhile goal for future work would be to prove that if all finite members of \mathcal{V} and \mathcal{W} are tractable, then the same will hold for all finite members of $\mathcal{V} \circ \mathcal{W}$. Tellingly, Freese and McKenzie show in [FM17] that a number of important properties are preserved by Mal’tsev product.

Theorem 1.2 ([FM17]). *Let \mathcal{V} and \mathcal{W} be idempotent varieties. For each of the following properties, P , if both \mathcal{V} and \mathcal{W} have P , then so does $\mathbf{H}(\mathcal{V} \circ \mathcal{W})$:*

- (1) *being idempotent;*
- (2) *having a Taylor term;*

- (3) *being SD- \wedge* ;
- (4) *having an edge term*.

It follows from Theorem 1.2 that if both \mathcal{V} and \mathcal{W} are SD- \wedge , or both have an edge term, then every finite member of $\mathbf{H}(\mathcal{V} \circ \mathcal{W})$ is tractable. So the next step would be to consider \mathcal{V} with one of these properties and \mathcal{W} with the other. But even if we could prove that tractability is preserved by Mal'tsev product, that would not be enough. In particular, simple algebras would remain problematic.

Interestingly, Barto and Kozik developed an approach that is particularly well-suited to dealing with simple algebras in this context. In LICS'10 ([BK10]), these researchers proved a powerful “Absorption Theorem” for products of two “absorption-free” algebras in a Taylor variety. At a more recent workshop [Bar15], Barto announced further joint work with Kozik on a general “Rectangularity Theorem” that gives conditions under which a subdirect product of simple nonabelian algebras contains a full product of “minimal absorbing” subalgebras. (Terms in quotes and other technicalities are defined below.)

In Section 5 of the present paper we state and prove a version of Barto and Kozik’s Rectangularity Theorem. In Section 8, we apply this tool, together with some techniques involving Mal'tsev products and other new decomposition strategies, to prove that every commutative, idempotent binar of cardinality at most 4 is tractable.

Note: Subsequent to the initial submission of this manuscript, Andrei Bulatov [Bul17] and Dmitriy Zhuk [Zhu17] independently announced a resolution of Conjecture 1.1. The proofs are ingenious, lengthy, and exceedingly complex, but overall seem to follow a strategy similar to the one outlined above. Nonetheless, we feel the methods and results presented here are useful for tackling specific problems, and doing so more directly than would be possible using the approaches offered by proofs of the general conjecture.

2. DEFINITIONS AND NOTATIONS

2.1. Notation for projections, scopes, and kernels. Here we collect definitions and notations used in the sequel. Some of these will be reintroduced more carefully later, as needed.

An operation $f: A^n \rightarrow A$ is called *idempotent* provided $f(a, a, \dots, a) = a$ for all $a \in A$. Examples of idempotent operations are the projection functions and these play an important role in later sections, so we start by introducing a flexible notation for them.

We define natural numbers as usual and denote them by $0 := \emptyset$, $\underline{n} := \{0, 1, 2, \dots, n-1\}$. Given sets A_0, A_1, \dots, A_{n-1} , An element $\mathbf{a} \in \prod_{i \in \underline{n}} A_i$ of their Cartesian product may be viewed as an ordered n -tuple by simply listing its values, $\mathbf{a} = (\mathbf{a}(0), \mathbf{a}(1), \dots, \mathbf{a}(n-1))$, or as a function defined on index set \underline{n} ; the latter view is emphasized symbolically as follows:

$$\mathbf{a}: \underline{n} \rightarrow \bigcup_{i \in \underline{n}} A_i; \quad i \mapsto \mathbf{a}(i) \in A_i.$$

If $\sigma: \underline{k} \rightarrow \underline{n}$ is a k -tuple of numbers in \underline{n} , then we can compose an n -tuple \mathbf{a} in $\prod_{i \in \underline{n}} A_i$ with σ yielding the k -tuple $\mathbf{a} \circ \sigma$ in $\prod_{i \in \underline{k}} A_{\sigma(i)}$. Generally speaking, we will try to avoid nonstandard notational conventions, but here are two exceptions:

$$\underline{A} := \prod_{i \in \underline{n}} A_i \quad \text{and} \quad \underline{A}_\sigma := \prod_{i \in \underline{k}} A_{\sigma(i)}.$$

Now, if the k -tuple $\sigma: \underline{k} \rightarrow \underline{n}$ happens to be one-one, and if we let p_σ denote the map $\mathbf{a} \mapsto \mathbf{a} \circ \sigma$, then p_σ is the usual *projection function* from \underline{A} onto \underline{A}_σ . Thus, $p_\sigma(\mathbf{a})$ is a k -tuple whose i -th component is $(\mathbf{a} \circ \sigma)(i) = \mathbf{a}(\sigma(i))$. We will make frequent use of such projections, as well as their images under the covariant and contravariant powerset functors \mathcal{P} and $\bar{\mathcal{P}}$. Indeed, we let $\text{Proj}_\sigma: \mathcal{P}(\underline{A}) \rightarrow \mathcal{P}(\underline{A}_\sigma)$ denote the *projection set function* defined for each $R \subseteq \underline{A}$ by $\text{Proj}_\sigma R = \mathcal{P}(p_\sigma)(R) = \{p_\sigma(\mathbf{x}) \mid \mathbf{x} \in R\} = \{\mathbf{x} \circ \sigma \mid \mathbf{x} \in R\}$, and we let $\text{Inj}_\sigma: \bar{\mathcal{P}}(\underline{A}_\sigma) \rightarrow \bar{\mathcal{P}}(\underline{A})$ denote the *injection set function* defined for each $S \subseteq \underline{A}_\sigma$ by

$$\text{Inj}_\sigma S = \bar{\mathcal{P}}(p_\sigma)(S) = \{\mathbf{x} \mid p_\sigma(\mathbf{x}) \in S\} = \{\mathbf{x} \in \underline{A} \mid \mathbf{x} \circ \sigma \in S\}. \quad (2.1)$$

Of course, $\text{Inj}_\sigma S$ is nothing more than the inverse image of the set S with respect to the projection function p_σ . We sometimes use the shorthand $R_\sigma = \text{Proj}_\sigma R$ and $S^{\bar{\sigma}} = \text{Inj}_\sigma S$ for the projection and injection set functions, respectively.

A one-one function $\sigma: \underline{k} \rightarrow \underline{n}$ is called a *scope* or *scope function*. If $R \subseteq \underline{A}_\sigma$ is a subset of the projection of \underline{A} onto coordinates $(\sigma(0), \sigma(1), \dots, \sigma(k-1))$, then we call R a *relation on \underline{A} with scope σ* . The pair (σ, R) is called a *constraint*, and $R^{\bar{\sigma}}$ is the set of tuples in \underline{A} that *satisfy* (σ, R) .

By η_σ we denote the *kernel* of the projection function p_σ , which is the following equivalence relation on \underline{A} :

$$\eta_\sigma = \{(\mathbf{a}, \mathbf{a}') \in \underline{A}^2 \mid p_\sigma(\mathbf{a}) = p_\sigma(\mathbf{a}')\} = \{(\mathbf{a}, \mathbf{a}') \in \underline{A}^2 \mid \mathbf{a} \circ \sigma = \mathbf{a}' \circ \sigma\}. \quad (2.2)$$

More generally, if θ is an equivalence relation on the set $\prod_{j \in \underline{k}} A_{\sigma(j)}$, then we define the equivalence relation θ_σ on the set \underline{A} by $\theta_\sigma = \{(\mathbf{a}, \mathbf{a}') \in \underline{A}^2 \mid (\mathbf{a} \circ \sigma) \theta (\mathbf{a}' \circ \sigma)\}$. Thus, θ_σ consists of pairs in \underline{A}^2 that land in θ when projected onto coordinates in the scope σ . If 0_A is the least equivalence on \underline{A} , then η_σ is shorthand for $(0_A)_\sigma$.

If the domain of σ is a singleton, $\underline{k} = \{0\}$, then σ is just a one-element list, say, $\sigma = (j)$, and we write p_j instead of $p_{(j)}$. Similarly, we write Proj_j instead of $\text{Proj}_{(j)}$, η_j instead of $\eta_{(j)}$, etc. Here are some obvious consequences of these definitions: $\bigvee_{j \in \underline{n}} \eta_j = \underline{A}^2$, $\eta_\sigma = \bigwedge_{j \in \sigma} \eta_j$, $\eta_{\underline{n}} = \bigwedge_{j \in \underline{n}} \eta_j = 0_A$, $\theta_\sigma = \bigwedge_{j \in \underline{k}} \theta_{\sigma(j)}$.

2.2. Notation for algebraic structures. This section introduces the notation we use for algebras and related concepts. The reader should consult [Ber12] for more details and background on general (universal) algebras and the varieties they inhabit.

2.2.1. Product algebras. Fix $n \in \mathbb{N}$, let F be a set of operation symbols, and for each $i \in \underline{n}$ let $\mathbf{A}_i = \langle A_i, F \rangle$ be an algebra of type F . Let $\mathbf{A} = \mathbf{A}_0 \times \mathbf{A}_1 \times \dots \times \mathbf{A}_{n-1} = \prod_{i \in \underline{n}} \mathbf{A}_i$ denote the product algebra. If $k \leq n$ and if $\sigma: \underline{k} \rightarrow \underline{n}$ is a one-one function, then we denote by $p_\sigma: \mathbf{A} \rightarrow \prod_{i \in \underline{k}} \mathbf{A}_{\sigma(i)}$ the *projection* of \mathbf{A} onto the “ σ -factors” of \mathbf{A} , which is an algebra homomorphism; thus the kernel η_σ defined in (2.2) is a congruence relation of \mathbf{A} .

2.2.2. Term operations. For a nonempty set A , we let \mathcal{O}_A denote the set of all finitary operations on A . That is, $\mathcal{O}_A = \bigcup_{n \in \mathbb{N}} A^{(A^n)}$. A *clone of operations* on A is a subset of \mathcal{O}_A that contains all projection operations and is closed under general composition. If $\mathbf{A} = \langle A, F^{\mathbf{A}} \rangle$ denotes the algebra with universe A and set of basic operations F , then $\text{Clo}(\mathbf{A})$ denotes the clone generated by F , which is also known as the *clone of term operations* of \mathbf{A} . We write $\text{Clo}_n(\mathbf{A})$ to denote the n -ary members of $\text{Clo}(\mathbf{A})$ (despite the fact that $\text{Clo}_n(\mathbf{A})$ is not a clone).

A *strong Mal'tsev condition* is a primitive positive sentence in the language of clones. A *Mal'tsev condition* is a countably infinite disjunction of strong Mal'tsev conditions $\{\varphi_i\}$ that get weaker as i increases (that is, $\varphi_i \vdash \varphi_{i+1}$). A *nontrivial Mal'tsev condition* is one that fails in some variety (equivalently, fails in the variety of sets). An *idempotent Mal'tsev condition* is one that satisfies the following: for each term f appearing in the condition, the identity $f(x, \dots, x) \approx x$ is a consequence of the identities in the Mal'tsev condition. (See, e.g., [Ber12] or [KK13] for more details.)

Walter Taylor proved in [Tay77] that a variety \mathcal{V} satisfies some nontrivial idempotent Mal'tsev condition if and only if it satisfies one of the following form: for some n , \mathcal{V} has an idempotent n -ary term t such that for each $i \in \underline{n}$ there is an identity of the form

$$t(*, \dots, *, x, *, \dots, *) \approx t(*, \dots, *, y, *, \dots, *)$$

true in \mathcal{V} where distinct variables x and y appear in the i -th position on either side of the identity. Such a term t is now commonly called a *Taylor term*.

Throughout this paper we assume all algebras are finite (though some results may apply more generally). Starting in Section 5, we make the additional assumption that the algebras in question come from a single *Taylor variety* \mathcal{V} , by which we mean that \mathcal{V} has a Taylor term and every term of \mathcal{V} is idempotent.

2.2.3. Subdirect products. If $k, n \in \mathbb{N}$, if $\mathcal{A} = (A_0, A_1, \dots, A_{n-1})$ is a list of sets, and if $\sigma: \underline{k} \rightarrow \underline{n}$ is a k -tuple of indices, then a relation R over \mathcal{A} with scope σ is a subset of the Cartesian product $A_{\sigma(0)} \times A_{\sigma(1)} \times \dots \times A_{\sigma(k-1)}$. Let F be a set of operation symbols and for each $i \in \underline{n}$ let $\mathbf{A}_i = \langle A_i, F \rangle$ be an algebra of type F . If $\mathbf{A} = \prod_{\underline{n}} \mathbf{A}_i$ is the product of these algebras, then a relation R over \mathcal{A} with scope σ is called *compatible with \mathbf{A}* if it is closed under the basic operations in F . In other words, R is compatible if the induced algebra $\mathbf{R} = \langle R, F \rangle$ is a subalgebra of $\prod_k \mathbf{A}_{\sigma(j)}$. If R is compatible with the product algebra and if the projection of R onto each factor is surjective, then \mathbf{R} is called a *subdirect product* of the algebras in the list $(\mathbf{A}_{\sigma(0)}, \mathbf{A}_{\sigma(1)}, \dots, \mathbf{A}_{\sigma(k-1)})$; we denote this situation by writing $\mathbf{R} \leq_{\text{sd}} \prod_k \mathbf{A}_{\sigma(j)}$.

3. ABELIAN ALGEBRAS

In later sections nonabelian algebras will play the following role: some of the theorems will begin with the assumption that a particular algebra \mathbf{A} is nonabelian and then proceed to show that if the result to be proved were false, then \mathbf{A} would have to be abelian. To prepare the way for such arguments, we review some basic facts about abelian algebras.

3.1. Definitions. Let $\mathbf{A} = \langle A, F^{\mathbf{A}} \rangle$ be an algebra. A reflexive, symmetric, compatible binary relation $T \subseteq A^2$ is called a *tolerance of \mathbf{A}* . Given a pair $(\mathbf{u}, \mathbf{v}) \in A^m \times A^m$ of m -tuples of A , we write $\mathbf{u} \mathbf{T} \mathbf{v}$ just in case $\mathbf{u}(i) T \mathbf{v}(i)$ for all $i \in \underline{m}$. We state a number of definitions in this section using tolerance relations, but the definitions don't change when the tolerance in question happens to be a congruence relation (i.e., a transitive tolerance).

Suppose S and T are tolerances on \mathbf{A} . An *S, T -matrix* is a 2×2 array of the form

$$\begin{bmatrix} t(\mathbf{a}, \mathbf{u}) & t(\mathbf{a}, \mathbf{v}) \\ t(\mathbf{b}, \mathbf{u}) & t(\mathbf{b}, \mathbf{v}) \end{bmatrix},$$

where t , \mathbf{a} , \mathbf{b} , \mathbf{u} , \mathbf{v} have the following properties:

- (i) $t \in \text{Clo}_{\ell+m}(\mathbf{A})$,
- (ii) $(\mathbf{a}, \mathbf{b}) \in A^\ell \times A^\ell$ and $\mathbf{a} \mathbf{S} \mathbf{b}$,
- (iii) $(\mathbf{u}, \mathbf{v}) \in A^m \times A^m$ and $\mathbf{u} \mathbf{T} \mathbf{v}$.

Let δ be a congruence relation of \mathbf{A} . If the entries of every S, T -matrix satisfy

$$t(\mathbf{a}, \mathbf{u}) \delta t(\mathbf{a}, \mathbf{v}) \iff t(\mathbf{b}, \mathbf{u}) \delta t(\mathbf{b}, \mathbf{v}), \quad (3.1)$$

then we say that S centralizes T modulo δ and we write $\mathbf{C}(S, T; \delta)$. That is, $\mathbf{C}(S, T; \delta)$ means that (3.1) holds for all $\ell, m, t, \mathbf{a}, \mathbf{b}, \mathbf{u}, \mathbf{v}$ satisfying properties (i)–(iii).

The commutator of S and T , denoted by $[S, T]$, is the least congruence δ such that $\mathbf{C}(S, T; \delta)$ holds. Note that $\mathbf{C}(S, T; 0_A)$ is equivalent to $[S, T] = 0_A$, and this is sometimes called the S, T -term condition; when it holds we say that S centralizes T , and write $\mathbf{C}(S, T)$. A tolerance T is called *abelian* if $\mathbf{C}(T, T)$ (i.e., $[T, T] = 0_A$). An algebra \mathbf{A} is called *abelian* if 1_A is abelian (i.e., $\mathbf{C}(1_A, 1_A)$).

3.2. Facts about centralizers and abelian congruences. We now collect some useful facts about centralizers of congruence relations that are needed in Section 5. The facts listed in the first lemma below are well-known and easy to prove. (For examples, see [HM88, Prop 3.4] and [KK13, Thm 2.19].)

Lemma 3.1. *Let \mathbf{A} be an algebra and suppose \mathbf{B} is a subalgebra of \mathbf{A} . Let $\alpha, \beta, \gamma, \delta, \alpha_i, \beta_j, \gamma_k$ be congruences of \mathbf{A} , for some $i \in I, j \in J, k \in K$. Then the following hold:*

- (1) $\mathbf{C}(\alpha, \beta; \alpha \wedge \beta)$;
- (2) if $\mathbf{C}(\alpha, \beta; \gamma_k)$ for all $k \in K$, then $\mathbf{C}(\alpha, \beta; \bigwedge_K \gamma_k)$;
- (3) if $\mathbf{C}(\alpha_i, \beta; \gamma)$ for all $i \in I$, then $\mathbf{C}(\bigvee_I \alpha_i, \beta; \gamma)$;
- (4) if $\mathbf{C}(\alpha, \beta; \gamma)$ and $\alpha' \leq \alpha$, then $\mathbf{C}(\alpha', \beta; \gamma)$;
- (5) if $\mathbf{C}(\alpha, \beta; \gamma)$ and $\beta' \leq \beta$, then $\mathbf{C}(\alpha, \beta'; \gamma)$;
- (6) if $\mathbf{C}(\alpha, \beta; \gamma)$ in \mathbf{A} , then $\mathbf{C}(\alpha \cap B^2, \beta \cap B^2; \gamma \cap B^2)$ in \mathbf{B} ;
- (7) if $\gamma \leq \delta$, then $\mathbf{C}(\alpha, \beta; \delta)$ in \mathbf{A} if and only if $\mathbf{C}(\alpha/\gamma, \beta/\gamma; \delta/\gamma)$ in \mathbf{A}/γ .

The next two lemmas are essential in a number proofs below. The first lemma identifies special conditions under which certain quotient congruences are abelian. The second gives fairly general conditions under which quotients of abelian congruences are abelian.

Lemma 3.2. *Let $\alpha_0, \alpha_1, \beta$ be congruences of \mathbf{A} and suppose $\alpha_0 \wedge \beta = \delta = \alpha_1 \wedge \beta$. Then $\mathbf{C}(\alpha_0 \vee \alpha_1, \beta; \delta)$. If, in addition, $\beta \leq \alpha_0 \vee \alpha_1$, then $\mathbf{C}(\beta, \beta; \delta)$, so β/δ is an abelian congruence of \mathbf{A}/δ .*

Lemma 3.2 is an easy consequence of items (1), (3), (4), and (7) of Lemma 3.1.

Lemma 3.3. *Let \mathcal{V} be a locally finite variety with a Taylor term and let $\mathbf{A} \in \mathcal{V}$. Then $\mathbf{C}(\beta, \beta; \gamma)$ for all $[\beta, \beta] \leq \gamma$.*

Lemma 3.3 can be proved by combining the next result, of David Hobby and Ralph McKenzie, with a result of Keith Kearnes and Emil Kiss.

Lemma 3.4 (cf. [HM88, Thm 7.12]). *A locally finite variety \mathcal{V} has a Taylor term if and only if it has a so called weak difference term; that is, a term $d(x, y, z)$ satisfying the following conditions for all $\mathbf{A} \in \mathcal{V}$, all $a, b \in A$, and all $\beta \in \text{Con}(\mathbf{A})$: $d^{\mathbf{A}}(a, a, b) [\beta, \beta] b [\beta, \beta]$ $d^{\mathbf{A}}(b, a, a)$, where $\beta = \text{Cg}^{\mathbf{A}}(a, b)$.*

Lemma 3.5 ([KK13, Lem 6.8]). *If \mathbf{A} belongs to a variety with a weak difference term and if β and γ are congruences of \mathbf{A} satisfying $[\beta, \beta] \leq \gamma$, then $\mathbf{C}(\beta, \beta; \gamma)$.*

Remark 3.6. It follows immediately from Lemma 3.3 that in a locally finite Taylor variety, \mathcal{V} , quotients of abelian algebras are abelian, so the abelian members of \mathcal{V} form a subvariety. But this can also be derived from Lemma 3.4, since $[\beta, \beta] = 0_A$ implies $d^{\mathbf{A}}$ is a Mal'tsev term operation on the blocks of β , so if \mathbf{A} is abelian—i.e., if $\mathbf{C}(1_A, 1_A; 0_A)$ —then Lemma 3.4, implies that \mathbf{A} has a Mal'tsev term operation. (This will be recorded below in Theorem 6.7.) It then follows that homomorphic images of \mathbf{A} are abelian. (See [Ber12, Cor 7.28] for more details).

4. ABSORPTION THEORY

In this section we survey some of the theory related to an important concept called *absorption*, which was invented by Libor Barto and Marcin Kozik. (In [BK17], these authors provide their own survey, to which we refer the reader for more details.) After introducing the concept, we discuss some of the properties that make absorption so useful. The main results in this section are not new. The only possibly novel contributions are some straightforward observations in Section 4.3 (though these have likely been observed by others). Our intention here is merely to collect and present these results in a way that makes them easy to apply in the sequel.

Let \mathbf{A} be an algebra and $t \in \text{Clo}(\mathbf{A})$ a k -ary term operation. A subalgebra $\mathbf{B} \leq \mathbf{A}$ is said to be *absorbing in \mathbf{A} with respect to t* if for all $1 \leq j \leq k$ and for all $(b_1, \dots, b_{j-1}, a, b_{j+1}, \dots, b_k) \in B^{j-1} \times A \times B^{k-j}$ we have $t(b_1, \dots, b_{j-1}, a, b_{j+1}, \dots, b_k) \in B$. In other terms, $t[B^{j-1} \times A \times B^{k-j}] \subseteq B$, for all $1 \leq j \leq k$, where $t[D]$ denotes the set $\{t(x) \mid x \in D\}$. The notation $\mathbf{B} \triangleleft \mathbf{A}$ and the phrase “ \mathbf{B} absorbs \mathbf{A} ” are used just in case \mathbf{B} is an absorbing subalgebra of \mathbf{A} with respect to some term. To be explicit about the term, one writes $\mathbf{B} \triangleleft_t \mathbf{A}$, and t is sometimes called an *absorbing term* for \mathbf{B} in this case. If B is a subuniverse of \mathbf{A} , then $B \triangleleft \mathbf{A}$ means that the subalgebra $\mathbf{B} = \langle B, \dots \rangle$, with universe B , absorbs \mathbf{A} . The notation $\mathbf{B} \triangleleft\triangleleft \mathbf{A}$ means that \mathbf{B} is a minimal absorbing subalgebra of \mathbf{A} , that is, B is minimal (with respect to set inclusion) among the absorbing subuniverses of \mathbf{A} . An algebra is said to be *absorption-free* if it has no proper absorbing subalgebras.

4.1. Absorption theorem. In later sections, we will make frequent use of a powerful tool of Barto and Kozik called the “Absorption Theorem” (Thm 2.3 of [BK12]). This result concerns the special class of “linked” subsets of products. A subset $R \subseteq A_0 \times A_1$ is said to be *linked* if it satisfies the following: for $a, a' \in \text{Proj}_0 R$ there exist elements $c_0, c_2, \dots, c_{2n} \in A_0$ and $c_1, c_3, \dots, c_{2n+1} \in A_1$ such that $c_0 = a$, $c_{2n} = a'$, and for all $0 \leq i < n$, $(c_{2i}, c_{2i+1}) \in R$ and $(c_{2i+2}, c_{2i+1}) \in R$. A subalgebra $\mathbf{R} \leq \mathbf{A}_0 \times \mathbf{A}_1$ linked just in case its universe is.

Here is an easily proved fact that provides some equivalent ways to define “linked.”

Fact 4.1. Let $\mathbf{R} = \langle R, \dots \rangle \leq_{\text{sd}} \mathbf{A}_0 \times \mathbf{A}_1$, let $\rho_i = \ker(R \rightarrow A_i)$ denote the kernel of the projection of R onto its i -th coordinate, and let $R^{-1} = \{(y, x) \in A_1 \times A_0 \mid (x, y) \in R\}$. Then the following are equivalent:

- (1) \mathbf{R} is linked;
- (2) $\rho_0 \vee \rho_1 = 1_R$;
- (3) if $a, a' \in \text{Proj}_0 R$, then (a, a') is in the transitive closure of $R \circ R^{-1}$.

Theorem 4.2 (Absorption Theorem [BK12, Thm 2.3]). *If \mathcal{V} is an idempotent locally finite variety, then the following are equivalent:*

- \mathcal{V} is a Taylor variety;
- if $\mathbf{A}_0, \mathbf{A}_1 \in \mathcal{V}$ are finite, absorption-free and $\mathbf{R} \leq_{\text{sd}} \mathbf{A}_0 \times \mathbf{A}_1$ is linked, then $\mathbf{R} = \mathbf{A}_0 \times \mathbf{A}_1$.

Before moving on to the next subsection, we need another definition. If $f: A^\ell \rightarrow A$ and $g: A^m \rightarrow A$ are ℓ -ary and m -ary operations on A , then by $f \star g$ we mean the ℓm -ary operation that maps $\mathbf{a} = (a_{11}, \dots, a_{1m}, a_{21}, \dots, a_{2m}, \dots, a_{\ell 1}, \dots, a_{\ell m}) \in A^{\ell m}$ to

$$(f \star g)(\mathbf{a}) = f(g(a_{11}, \dots, a_{1m}), g(a_{21}, \dots, a_{2m}), \dots, g(a_{\ell 1}, \dots, a_{\ell m})).$$

4.2. Properties of absorption. In this section we prove some easy facts about algebras that hold quite generally. In particular, we don't assume the presence of Taylor terms or other Mal'tsev conditions in this section. However, for the sake of simplicity, we do assume that all algebras are finite and belong to a single idempotent variety, though some of the results of this section hold for infinite algebras as well. Almost all of the results in this section are well known. The only exceptions are mild generalizations of facts that have already appeared in print. We have restated or reworked some of the known results in a way that makes them easier for us to apply.

Lemma 4.3. *Let \mathbf{A} be an algebra with term operations f and g . If $\mathbf{B} \triangleleft \mathbf{A}$ with respect to either f or g , then $\mathbf{B} \triangleleft \mathbf{A}$ with respect to $f \star g$.*

The straightforward proof appears in Section A.1. The next corollary is easy to prove by induction on p .

Corollary 4.4. *Fix $p < \omega$, a variety \mathcal{V} , algebras $\mathbf{A}, \mathbf{B} \in \mathcal{V}$, and terms t_0, t_1, \dots, t_{p-1} in the language of \mathcal{V} . If $\mathbf{B} \triangleleft_{t_i} \mathbf{A}$ for some $0 \leq i < p$, then $\mathbf{B} \triangleleft_s \mathbf{A}$, where $s = t_0 \star t_1 \star \dots \star t_{p-1}$.*

Lemma 4.5 ([BK12, Prop 2.4]). *Let \mathbf{A} be an algebra.*

- If $\mathbf{C} \triangleleft \mathbf{B} \triangleleft \mathbf{A}$, then $\mathbf{C} \triangleleft \mathbf{A}$.
- If $\mathbf{B} \triangleleft_f \mathbf{A}$, $\mathbf{C} \triangleleft_g \mathbf{A}$, $B \cap C \neq \emptyset$, and $t = f \star g$, then $B \cap C \triangleleft_t \mathbf{A}$.

The following is an easy corollary of Lemma 4.5 (for proof see Sec. A.2) and so is the generalization to finite products and minimal absorbing subalgebras given in Lemma 4.7 (and proved in Sec. A.3).

Corollary 4.6. *Suppose $\mathbf{B}_i \triangleleft_{f_i} \mathbf{A}_i$ for $i \in \{0, 1\}$ and $t = f_0 \star f_1$. Then $\mathbf{B}_0 \times \mathbf{B}_1 \triangleleft_t \mathbf{A}_0 \times \mathbf{A}_1$.*

Lemma 4.7. *Let $\mathbf{B}_i \leq \mathbf{A}_i$ ($0 \leq i < n$) be algebras and define $\mathbf{B} := \prod_n \mathbf{B}_i$ and $\mathbf{A} := \prod_n \mathbf{A}_i$. If $\mathbf{B}_i \triangleleft_{f_i} \mathbf{A}_i$ (resp., $\mathbf{B}_i \triangleleft_{\triangleleft_{f_i}} \mathbf{A}_i$) for each $0 \leq i < n$, then $\mathbf{B} \triangleleft_t \mathbf{A}$ (resp., $\mathbf{B} \triangleleft_{\triangleleft_t} \mathbf{A}$) where $t := f_0 \star f_1 \star \dots \star f_{n-1}$.*

Here is another simple observation with a trivial proof that is nonetheless quite useful for proving certain algebras absorption-free.

Lemma 4.8. *Let \mathbf{B} and \mathbf{C} be subalgebras of a finite idempotent algebra \mathbf{A} . Suppose $\mathbf{B} \triangleleft_t \mathbf{A}$, and suppose $D = B \cap C \neq \emptyset$. Then the restriction of t to C is an absorbing term for D in C , whence $D \triangleleft C$.*

Lemma 4.7 and Lemma 4.8 yield the following observation that we will need in Section 6:

Corollary 4.9. *Let $\mathbf{A}_0, \mathbf{A}_1, \dots, \mathbf{A}_{n-1}$ be finite algebras with $\mathbf{B}_i \triangleleft \mathbf{A}_i$ for each i . Suppose $\mathbf{R} \leq \prod_i \mathbf{A}_i$ and $R \cap \prod_i B_i \neq \emptyset$. Then $R \cap \prod_i B_i$ is an absorbing subuniverse of \mathbf{R} .*

For $0 < j \leq k$, we say that a k -ary term operation t of an algebra \mathbf{A} depends on its j -th argument provided there exist a_0, a_1, \dots, a_{k-1} such that $p(x) := t(a_0, \dots, a_{j-2}, x, a_j, \dots, a_{k-1})$ is a nonconstant polynomial of \mathbf{A} .

Let \mathbf{A} be an algebra, and $s \in C \subseteq A$. We call s a *sink* for C provided the following holds: for every term $t \in \text{Clo}_k(\mathbf{A})$ and all $0 < j \leq k$, if t depends on its j -th argument, then $t(c_0, c_1, \dots, c_{j-2}, s, c_j, \dots, c_{k-1}) = s$ for all $c_i \in C$. If B is an absorbing subuniverse that intersects nontrivially with a set containing a sink, then B must also contain the sink.

Lemma 4.10. *If $B \triangleleft A$, if s is a sink for $C \subseteq A$, and if $B \cap C \neq \emptyset$, then $s \in B$.*

We conclude this subsection with three more properties of absorption that will be useful below. The first of these is proved in Section A.5 below.

Lemma 4.11. *Let $\mathbf{A}_0, \mathbf{A}_1, \dots, \mathbf{A}_{n-1}$ be finite idempotent algebras of the same type, and suppose $\mathbf{B}_i \triangleleft \mathbf{A}_i$ for $0 \leq i < n$. Let $\mathbf{R} \leq_{\text{sd}} \mathbf{A}_0 \times \mathbf{A}_1 \times \dots \times \mathbf{A}_{n-1}$, and let $R' = R \cap (B_0 \times B_1 \times \dots \times B_{n-1})$. If $R' \neq \emptyset$, then $R' \leq_{\text{sd}} \mathbf{B}_0 \times \mathbf{B}_1 \times \dots \times \mathbf{B}_{n-1}$.*

Lemma 4.12 ([BK12, cf. Prop 2.15]). *Let \mathbf{A}_0 and \mathbf{A}_1 be finite idempotent algebras of the same type, let $\mathbf{R} \leq_{\text{sd}} \mathbf{A}_0 \times \mathbf{A}_1$ and assume \mathbf{R} is linked. If $\mathbf{S} \triangleleft \mathbf{R}$, then \mathbf{S} is also linked.*

Lemma 4.13 ([BKS15, Lem 4.1]). *Finite idempotent abelian algebras are absorption-free.*

4.3. Linking is easy, sometimes. We will apply the Absorption Theorem frequently below, so we pause here to consider one of the hypotheses of the theorem that might seem less familiar to some readers. Specifically, we consider when we should expect a subdirect product to be *linked*, as required of $\mathbf{R} \leq_{\text{sd}} \mathbf{A}_0 \times \mathbf{A}_1$ in the statement of the Absorption Theorem. We show that in some special cases the hypothesis comes essentially for free, courtesy of the following elementary property of subdirect products:

Lemma 4.14. *Let \mathbf{A}_0 and \mathbf{A}_1 be algebras. Suppose $\mathbf{R} \leq_{\text{sd}} \mathbf{A}_0 \times \mathbf{A}_1$ and let $\rho_i = \ker(\mathbf{R} \rightarrow \mathbf{A}_i)$ denote the kernel of the i -th projection of \mathbf{R} onto \mathbf{A}_i .*

- (1) *If \mathbf{A}_0 is simple, then either $\rho_0 \vee \rho_1 = 1_R$ or $\rho_0 \geq \rho_1$.*
- (2) *If \mathbf{A}_0 and \mathbf{A}_1 are both simple, then either $\rho_0 \vee \rho_1 = 1_R$ or $\rho_0 = 0_R = \rho_1$.*

Proof. $\mathbf{R}/\rho_0 \cong \mathbf{A}_0$ and \mathbf{A}_0 is simple, so ρ_0 is a coatom (i.e., maximal) in the congruence lattice $\text{Con } \mathbf{R}$. Thus, $\rho_0 \vee \rho_1 < 1_R$ would imply $\rho_0 = \rho_0 \vee \rho_1$, proving (1). If \mathbf{A}_1 is also simple, then the same argument, with the roles of ρ_0 and ρ_1 swapped, shows that $\rho_0 \leq \rho_1$, so $\rho_0 = \rho_1$. Since $\rho_0 \wedge \rho_1 = 0_R$, both projection kernels are 0_R in this case. \square

Corollary 4.15. *Suppose $\mathbf{A}_0, \mathbf{A}_1$ are simple and $\mathbf{R} \leq_{\text{sd}} \mathbf{A}_0 \times \mathbf{A}_1$. If $\rho_0 \neq \rho_1$ or $\mathbf{A}_0 \not\cong \mathbf{A}_1$, then \mathbf{R} is linked.*

In a locally finite Taylor variety, when one factor of $\mathbf{A}_0 \times \mathbf{A}_1$ is abelian and the other is simple and nonabelian, linking comes for free; that is, every subdirect product of such $\mathbf{A}_0 \times \mathbf{A}_1$ is linked, as we now prove.

Corollary 4.16. *Suppose \mathbf{A}_0 and \mathbf{A}_1 are algebras in a locally finite Taylor variety. If \mathbf{A}_0 is abelian and \mathbf{A}_1 is simple and nonabelian, then every subdirect product of $\mathbf{A}_0 \times \mathbf{A}_1$ is linked. Moreover, if $\mathbf{R} \leq_{\text{sd}} \mathbf{A}_0 \times \mathbf{A}_1$ and if $\mathbf{B}_1 \triangleleft \mathbf{A}_1$, then \mathbf{R} intersects $\mathbf{A}_0 \times \mathbf{B}_1$ nontrivially and this intersection forms a linked subdirect product of $\mathbf{A}_0 \times \mathbf{B}_1$.*

Proof. Suppose $\mathbf{R} \leq_{\text{sd}} \mathbf{A}_0 \times \mathbf{A}_1$ is not linked. Since \mathbf{A}_1 is simple, $\rho_0 \leq \rho_1$ by Lemma 4.14. Thus, $\rho_0 = \rho_0 \wedge \rho_1 = 0_R$, so $\mathbf{R} \cong \mathbf{R}/\rho_0 \cong \mathbf{A}_0$ is abelian, while $\mathbf{R}/\rho_1 \cong \mathbf{A}_1$ is not. This is a contradiction since \mathbf{R}/θ is abelian for all $\theta \in \text{Con } \mathbf{R}$ (by Remark 3.6). For the second part, since \mathbf{R} is subdirect, it follows that $R \cap (A_0 \times B_1)$ is nonempty and subdirect by Lemmas 4.12 and 4.13. By Corollary 4.9, $\mathbf{R} \cap (\mathbf{A}_0 \times \mathbf{B}_1) \triangleleft \mathbf{R}$, so Lemma 4.12 implies that the intersection is also linked. \square

We can extend the previous result to multiple abelian factors by collecting such factors into one. Let $\underline{n} := \{0, 1, \dots, n-1\}$ and $k' := \underline{n} - \{k\}$ (in particular, $0' := \{1, 2, \dots, n-1\}$), and let $\mathbf{R}_{0'} := \text{Proj}_{0'} \mathbf{R}$.

Corollary 4.17. *Let $\mathbf{A}_0, \mathbf{A}_1, \dots, \mathbf{A}_{n-1}$ be algebras in a locally finite Taylor variety. Suppose \mathbf{A}_0 is simple, nonabelian, and $\mathbf{A}_1, \mathbf{A}_2, \dots, \mathbf{A}_{n-1}$ are abelian. If $\mathbf{R} \leq_{\text{sd}} \prod_{\underline{n}} \mathbf{A}_i$, then $\mathbf{R} \leq_{\text{sd}} \mathbf{A}_0 \times \mathbf{R}_{0'}$ is linked.*

Proof. Suppose $\mathbf{R} \leq_{\text{sd}} \mathbf{A}_0 \times \mathbf{R}_{0'}$ is not linked: $\rho_0 \vee \rho_{0'} < 1_R$. Since \mathbf{A}_0 is simple, ρ_0 is a coatom, so $\rho_0 \geq \rho_{0'}$. Therefore, $\rho_{0'} = \rho_0 \wedge \rho_{0'} = 0_R$, so $\mathbf{R} \cong \mathbf{R}/\rho_{0'} \cong \mathbf{R}_{0'} \leq \prod_{0'} \mathbf{A}_i$. This proves that \mathbf{R} is abelian, yet $\mathbf{R}/\rho_0 \cong \mathbf{A}_0$ is not, which contradicts Remark 3.6. \square

The Absorption Theorem, Lemmas 4.11–4.13, and Corollaries 4.9, 4.15, 4.16 yield

Lemma 4.18. *Let \mathbf{A}_0 and \mathbf{A}_1 be finite algebras in a Taylor variety with minimal absorbing subalgebras $\mathbf{B}_i \triangleleft \mathbf{A}_i$ ($i = 0, 1$), and suppose $\mathbf{R} \leq_{\text{sd}} \mathbf{A}_0 \times \mathbf{A}_1$, where $\rho_0 \neq \rho_1$.*

- (i) *If \mathbf{A}_0 and \mathbf{A}_1 are simple and $R \cap (B_0 \times B_1) \neq \emptyset$, then $\mathbf{B}_0 \times \mathbf{B}_1 \leq \mathbf{R}$.*
- (ii) *If \mathbf{A}_0 is simple nonabelian and \mathbf{A}_1 is abelian, then $\mathbf{B}_0 \times \mathbf{A}_1 \leq \mathbf{R}$.*

Proof.

- (i) First note that \mathbf{R} is linked, by Corollary 4.15. Let $\mathbf{R}' := \mathbf{R} \cap (\mathbf{B}_0 \times \mathbf{B}_1)$. Then by Lemma 4.11 $\mathbf{R}' \leq_{\text{sd}} \mathbf{B}_0 \times \mathbf{B}_1$, and by Corollary 4.9 $\mathbf{R}' \triangleleft \mathbf{R}$, so \mathbf{R}' is also linked, by Lemma 4.12. The hypotheses of the Absorption Theorem—with \mathbf{R}' in place of \mathbf{R} and \mathbf{B}_i in place of \mathbf{A}_i —are now satisfied. Therefore, $\mathbf{R}' = \mathbf{B}_0 \times \mathbf{B}_1$.
- (ii) This follows directly from Corollary 4.16, Lemma 4.13, and the Absorption Theorem. \square

Again, collecting multiple abelian factors into one, we obtain

Corollary 4.19. *Let $\mathbf{A}_0, \mathbf{A}_1, \dots, \mathbf{A}_{n-1}$ be finite Taylor algebras, where \mathbf{A}_0 is simple nonabelian, $\mathbf{B}_0 \triangleleft \mathbf{A}_0$, and the remaining \mathbf{A}_i are abelian, and suppose $\mathbf{R} \leq_{\text{sd}} \prod_{\underline{n}} \mathbf{A}_i$. Then $\mathbf{B}_0 \times \mathbf{R}_{0'} \leq \mathbf{R}$.*

Proof. Obviously, $\mathbf{R} \leq_{\text{sd}} \mathbf{A}_0 \times \mathbf{R}_{0'}$ and $R \cap (B_0 \times R_{0'}) \neq \emptyset$. Also, $\mathbf{R}_{0'} := \text{Proj}_{0'} \mathbf{R}$ is abelian, so we can apply Lemma 4.18 (ii), with $\mathbf{R}_{0'}$ in place of \mathbf{A}_1 . \square

5. THE RECTANGULARITY THEOREM

5.1. Some history. In late May of 2015 we attended a workshop on “Open Problems in Universal Algebra,” at which Libor Barto announced a new theorem that he and Marcin Kozik proved called the “Rectangularity Theorem.” At this meeting Barto gave a detailed overview of the proof [Bar15]. The authors of the present paper then made a concerted effort over many months to fill in the details and produce a complete proof. Unfortunately,

each attempt uncovered a gap that we were unable to fill. Nonetheless, we found a slightly different route to the theorem. Our argument is similar to the one presented by Barto in most of its key aspects. In particular, we make heavy use of absorption and the Absorption Theorem plays a key role. However, we were unable to complete the proof without a new “Linking Lemma” (Lem. 5.4) that we proved using a general result of Kearnes and Kiss. Thus, our argument is similar in spirit to the original, but is (at the time of this writing, to our knowledge) the first complete proof of Barto and Kozik’s Rectangularity Theorem to appear in print. We hope our presentation sheds additional light on the result.

In the next subsection, we prove the Rectangularity Theorem and some corollaries that we will apply in Section 6, where we show how to use these results to determine the complexity of certain constraint satisfaction problems.

5.2. Preliminaries. We are mainly concerned with finite algebras belonging to a single Taylor variety \mathcal{V} (though some results below hold more generally). As mentioned, our goal in this subsection is a proof of the Rectangularity Theorem of Barto and Kozik. The theorem can be described informally as follows: Start with finite algebras $\mathbf{A}_0, \mathbf{A}_1, \dots, \mathbf{A}_{n-1}$ in a Taylor variety with minimal absorbing subalgebras $\mathbf{B}_i \triangleleft\triangleleft \mathbf{A}_i$. Suppose at most one of these algebras is abelian and the nonabelian algebras are simple. Let \mathbf{R} be a subdirect product of $\prod_{\underline{n}} \mathbf{A}_i$. Then (under a few more assumptions) if R intersects nontrivially with the product $\prod_{\underline{n}} B_i$, then R contains the entire product $\prod_{\underline{n}} B_i$.

Our proof of this result will depend on a number of useful lemmas that we now prove.

Notation.

- $\underline{n} := \{0, 1, 2, \dots, n-1\}$.

- For $\sigma \subseteq \underline{n}$, $A = A_0 \times A_1 \times \dots \times A_{n-1}$, and $\mathbf{a} = (a_0, a_1, \dots, a_{n-1})$, let

$$\eta_\sigma := \ker(A \rightarrow \prod_{\sigma} A_i) = \{(\mathbf{a}, \mathbf{a}') \in A^2 \mid a_i = a'_i \text{ for all } i \in \sigma\}$$

be the kernel of the projection of A onto coordinates σ .

- For $\mathbf{R} \leq_{\text{sd}} \mathbf{A}_0 \times \mathbf{A}_1 \times \dots \times \mathbf{A}_{n-1}$, let

$$\rho_\sigma := \ker(R \rightarrow \prod_{\sigma} A_i) = \{(\mathbf{r}, \mathbf{r}') \in R^2 \mid r_i = r'_i \text{ for all } i \in \sigma\}$$

be the kernel of the projection of R onto the coordinates σ . Thus, $\rho_\sigma = \eta_\sigma \cap R^2$.

- For $\sigma \subseteq \underline{n}$, let $\sigma' := \underline{n} - \sigma$, and take $\mathbf{R} \leq_{\text{sd}} \prod_{\sigma} \mathbf{A}_i \times \prod_{\sigma'} \mathbf{A}_i$ to mean the following:

(1) \mathbf{R} is a subalgebra of $\prod_{\underline{n}} \mathbf{A}_i$;

(2) $\text{Proj}_\sigma \mathbf{R} = \prod_{\sigma} \mathbf{A}_i$;

(3) $\text{Proj}_{\sigma'} \mathbf{R} = \prod_{\sigma'} \mathbf{A}_i$;

we say that \mathbf{R} is a *subdirect product* of $\prod_{\sigma} \mathbf{A}_i$ and $\prod_{\sigma'} \mathbf{A}_i$ in this case.

- The subdirect product $\mathbf{R} \leq_{\text{sd}} \prod_{\sigma} \mathbf{A}_i \times \prod_{\sigma'} \mathbf{A}_i$ is said to be *linked* if $\rho_\sigma \vee \rho_{\sigma'} = 1_R$.
- We sometimes use \mathbf{R}_σ as shorthand for $\text{Proj}_\sigma \mathbf{R}$, the projection of \mathbf{R} onto coordinates σ .

5.3. Rectangularity theorem. With a few more preliminary results, we will be ready to prove the general Rectangularity Theorem. The first of these comes from combining Lemma 4.7, transitivity of absorption, Corollary 4.9, and Lemma 4.12.

Lemma 5.1. *Let $\mathbf{A}_0, \mathbf{A}_1, \dots, \mathbf{A}_{n-1}$ be finite algebras in a Taylor variety, let $\mathbf{B}_i \triangleleft\triangleleft \mathbf{A}_i$ for each $i \in \underline{n}$, and let $\underline{n} = \sigma \cup \sigma'$ be a partition of $\{0, 1, \dots, n-1\}$ into two nonempty disjoint subsets. Assume \mathbf{R} is a linked subdirect product of $\prod_{\sigma} \mathbf{A}_i$ and $\prod_{\sigma'} \mathbf{A}_i$, and suppose $R' := R \cap \prod_{\underline{n}} B_i \neq \emptyset$. Then $\prod_{\underline{n}} B_i \subseteq R$, so $\mathbf{R}' = \prod_{\underline{n}} \mathbf{B}_i$.*

Proof. By Lemma 4.11, $\mathbf{R}' \leq_{\text{sd}} \prod_{\sigma} \mathbf{B}_i \times \prod_{\sigma'} \mathbf{B}_i$, and by Corollary 4.9, $\mathbf{R}' \triangleleft \mathbf{R}$. Therefore, Lemma 4.12 implies \mathbf{R}' is linked. By Lemma 4.7 and transitivity of absorption, it follows that $\prod_{\sigma} \mathbf{B}_i$ and $\prod_{\sigma'} \mathbf{B}_i$ are both absorption-free, so the Absorption Theorem implies that $\mathbf{R}' = \prod_{\sigma} \mathbf{B}_i \times \prod_{\sigma'} \mathbf{B}_i$. \square

Next, we recall a theorem of Keith Kearnes and Emil Kiss, proved in [KK13], about algebras in a Taylor variety.

Theorem 5.2 ([KK13, Thm 3.27]). *Suppose α and β are congruences of a Taylor algebra. Then $\mathbf{C}(\alpha, \alpha; \alpha \wedge \beta)$ if and only if $\mathbf{C}(\alpha \vee \beta, \alpha \vee \beta; \beta)$.*

The context in which we apply Theorem 5.2 is that of the following corollary.

Corollary 5.3. *Let $\alpha_0, \alpha_1, \beta, \delta$ be congruences of a Taylor algebra \mathbf{A} , and suppose that $\alpha_0 \wedge \beta = \delta = \alpha_1 \wedge \beta$ and $\alpha_0 \vee \alpha_1 = \alpha_0 \vee \beta = \alpha_1 \vee \beta = 1_A$. Then, $\mathbf{C}(1_A, 1_A; \alpha_0)$ and $\mathbf{C}(1_A, 1_A; \alpha_1)$, so \mathbf{A}/α_0 and \mathbf{A}/α_1 are abelian algebras.*

Proof. The hypotheses of Lemma 3.2 hold, so $\mathbf{C}(\beta, \beta; \delta)$ and β/δ is an abelian congruence of \mathbf{A}/δ . Now, since $\delta = \alpha_i \wedge \beta$, we have $\mathbf{C}(\beta, \beta; \alpha_i \wedge \beta)$, so Theorem 5.2 implies $\mathbf{C}(\alpha_i \vee \beta, \alpha_i \vee \beta; \alpha_i)$. This yields $\mathbf{C}(1_A, 1_A; \alpha_i)$, since $\alpha_i \vee \beta = 1_A$. By Lemma 3.1 (7) then, $\mathbf{C}(1_A/\alpha_i, 1_A/\alpha_i; \alpha_i/\alpha_i)$. Equivalently, $\mathbf{C}(1_{A/\alpha_i}, 1_{A/\alpha_i}; 0_{A/\alpha_i})$. That is, \mathbf{A}/α_i is abelian. \square

Before proceeding, let us remark that the statements of many results to follow will include the hypothesis that subdirect products do not involve redundant factors (since such factors can be combined into one without changing the solution set of the corresponding CSP instance). This is expressed by assuming that kernels of projections onto distinct coordinates are distinct, that is, $\rho_i \neq \rho_j$ for all $i \neq j$.

Lemma 5.4 (Linking Lemma). *Let $n \geq 2$ and $k' := n - \{k\}$. Let $\mathbf{A}_0, \mathbf{A}_1, \dots, \mathbf{A}_{n-1}$ be finite algebras in a Taylor variety, and let $\mathbf{B}_i \triangleleft \mathbf{A}_i$ for all $0 \leq i < n$. Suppose*

- at most one \mathbf{A}_i is abelian,
- all nonabelian factors are simple,
- $\mathbf{R} \leq_{\text{sd}} \mathbf{A}_0 \times \mathbf{A}_1 \times \dots \times \mathbf{A}_{n-1}$, and
- $\rho_i \neq \rho_j$ for all $i \neq j$.

Then there exists k such that $\mathbf{R} \leq_{\text{sd}} \mathbf{A}_k \times \mathbf{R}_{k'}$ is linked.

Proof. By way of contradiction, suppose $\rho_k \vee \rho_{k'} < 1_R$ for all $0 \leq k < n$.

Case 1. Every \mathbf{A}_i is nonabelian.

If $n = 2$, the result holds by Corollary 4.15. Assume $n > 2$. Since every factor is simple, each ρ_k is a coatom, so $\rho_k \vee \rho_{k'} < 1_R$ implies $\rho_k \geq \rho_{k'}$. Therefore,

$$\rho_{k'} := \bigwedge_{i \neq k} \rho_i = \left(\bigwedge_{i \neq k} \rho_i \right) \wedge \rho_k = \rho_n = 0_R. \quad (5.1)$$

Note that (5.1) holds for all $k \in n$. Now, let $\tau \subseteq n$ be a subset that is maximal among those satisfying $\rho_\tau > 0_R$. By what we just showed, τ omits at least two indices, say, j and ℓ . By maximality, $\rho_\tau \wedge \rho_j = 0_R = \rho_\tau \wedge \rho_\ell$, and $\rho_\tau \not\leq \rho_j$, and $\rho_\tau \not\leq \rho_\ell$. Since all factors are simple, ρ_j and ρ_ℓ are (distinct) coatoms, so $\rho_\tau \vee \rho_j = \rho_\tau \vee \rho_\ell = \rho_j \vee \rho_\ell = 1_R$. Therefore, by Corollary 5.3, both $\mathbf{A}_j \cong \mathbf{R}/\rho_j$ and $\mathbf{A}_\ell \cong \mathbf{R}/\rho_\ell$ are abelian, which contradicts the assumption that there are no abelian factors.

Case 2. Exactly one \mathbf{A}_i is abelian.

Assume without loss of generality that \mathbf{A}_0 is the abelian factor. Then, for $k > 0$, \mathbf{A}_k is simple and nonabelian, so ρ_k is a coatom of $\text{Con } \mathbf{R}$. Therefore, $\rho_k \vee \rho_{k'} < 1_R$ implies $\rho_{k'} \leq \rho_k$, so

$$\rho_{k'} = \rho_{k'} \wedge \rho_k = \bigwedge_{i \in \underline{n}} \rho_i = 0_R. \quad (5.2)$$

(Note that this holds for every $0 < k < n$.)

Now, let $\theta = \rho_{0'} := \bigwedge_{i \neq 0} \rho_i$. Let τ be a maximal subset of $0'$ such that $\rho_\tau > \theta$. Obviously τ is a proper subset of $0'$. Moreover, by (5.2) there exists $j \in 0'$ such that $\rho_\tau \not\leq \rho_j$. Since ρ_j is a coatom, $\rho_\tau \vee \rho_j = 1_R$, and by maximality of τ , we have $\rho_\tau \wedge \rho_j = \theta$.

Theorem 5.2 yields $\mathbf{C}(\rho_\tau \vee \rho_j, \rho_\tau \vee \rho_j; \rho_j)$ if and only if $\mathbf{C}(\rho_\tau, \rho_\tau; \rho_\tau \wedge \rho_j)$; equivalently,

$$\mathbf{C}(1_R, 1_R; \rho_j) \iff \mathbf{C}(\rho_\tau, \rho_\tau; \theta).$$

However, $\mathbf{C}(1_R, 1_R; \rho_j)$ implies \mathbf{A}_j is abelian, but \mathbf{A}_0 is the only abelian factor, so

$$\mathbf{C}(\rho_\tau, \rho_\tau; \theta) \text{ does not hold.} \quad (5.3)$$

Since \mathbf{A}_0 is abelian, $\mathbf{C}(1_R, 1_R; \rho_0)$; *a fortiori*, $\mathbf{C}(\rho_\tau \vee \rho_0, \rho_\tau \vee \rho_0; \rho_0)$, which holds if and only if $\mathbf{C}(\rho_\tau, \rho_\tau; \rho_\tau \wedge \rho_0)$ by Theorem 5.2. Now, if $\rho_0 \wedge \rho_\tau \leq \theta$, then Lemma 3.3 implies $\mathbf{C}(\rho_\tau, \rho_\tau; \theta)$, which is false by (5.3). Thus, $\rho_0 \wedge \rho_\tau \not\leq \theta$. It follows that $\rho_0 \wedge \rho_\tau \neq 0_R$. By (5.2), then, there are at least two distinct indices, say, j and ℓ , in \underline{n} that do not belong to τ . By maximality of τ , $\rho_\tau \wedge \rho_j = \theta = \rho_\tau \wedge \rho_\ell$. By Corollary 5.3 (with $\alpha_0 = \rho_j$, $\alpha_1 = \rho_\ell$, $\beta = \rho_\tau$, $\delta = \theta$), it follows that $\mathbf{R}/\rho_j \cong \mathbf{A}_j$ and $\mathbf{R}/\rho_\ell \cong \mathbf{A}_\ell$ are both abelian—a contradiction. \square

We have assembled all the tools we will use to accomplish the main goal of this subsection, which is to prove the following:

Theorem 5.5 (Rectangularity Theorem, cf. [BK17, Theorem 38]). *Let $\mathbf{A}_0, \mathbf{A}_1, \dots, \mathbf{A}_{n-1}$ be finite algebras in a Taylor variety, and let $\mathbf{B}_i \triangleleft \mathbf{A}_i$ for all $0 \leq i < n$. Suppose*

- *at most one \mathbf{A}_i is abelian,*
- *all nonabelian factors are simple,*
- $\mathbf{R} \leq_{\text{sd}} \mathbf{A}_0 \times \mathbf{A}_1 \times \dots \times \mathbf{A}_{n-1}$,
- $\rho_i \neq \rho_j$ for all $i \neq j$,
- $R' := R \cap (B_0 \times B_1 \times \dots \times B_{n-1}) \neq \emptyset$.

Then, $\mathbf{R}' = \mathbf{B}_0 \times \mathbf{B}_1 \times \dots \times \mathbf{B}_{n-1}$.

Proof. It suffices to prove $B_0 \times B_1 \times \dots \times B_{n-1} \subseteq R$, which we do by induction on the number of factors in the product $\mathbf{A}_0 \times \mathbf{A}_1 \times \dots \times \mathbf{A}_{n-1}$.

For $n = 2$ the result holds by Lemma 4.18. Fix $n > 2$ and assume for all $2 \leq k < n$ that the result holds for subdirect products of k factors. We prove the result for subdirect products of n factors.

Let σ be a nonempty proper subset of \underline{n} (so $1 \leq |\sigma| < n$). As above σ' is the complement of σ in \underline{n} , and \mathbf{R}_σ is the projection of \mathbf{R} onto coordinates in σ . Each of the projections \mathbf{R}_σ and $\mathbf{R}_{\sigma'}$ satisfies the assumptions of the theorem. (Indeed, $\mathbf{R}_\sigma \leq_{\text{sd}} \prod_{\sigma} \mathbf{A}_i$ and for all $i \neq j$ in σ , $\ker(R_\sigma \twoheadrightarrow A_i) \neq \ker(R_\sigma \twoheadrightarrow A_j)$; similarly for $\mathbf{R}_{\sigma'}$.) Therefore, the induction hypothesis implies $\prod_{\sigma} \mathbf{B}_i \leq \mathbf{R}_\sigma$ and $\prod_{\sigma'} \mathbf{B}_i \leq \mathbf{R}_{\sigma'}$. By Lemma 4.7, $\prod_{\sigma} \mathbf{B}_i \triangleleft \prod_{\sigma} \mathbf{A}_i$, and from $\prod_{\sigma} \mathbf{B}_i \leq \mathbf{R}_\sigma \leq \prod_{\sigma} \mathbf{A}_i$ follows $\prod_{\sigma} \mathbf{B}_i \triangleleft \mathbf{R}_\sigma$. In fact, $\prod_{\sigma} \mathbf{B}_i \triangleleft \mathbf{R}_\sigma$ as well, by minimality of $\prod_{\sigma} \mathbf{B}_i \triangleleft \prod_{\sigma} \mathbf{A}_i$, and transitivity of absorption. To summarize, for every $\emptyset \subsetneq \sigma \subsetneq \underline{n}$,

$$\mathbf{R} \leq_{\text{sd}} \mathbf{R}_\sigma \times \mathbf{R}_{\sigma'}, \quad \prod_{\sigma} \mathbf{B}_i \triangleleft \mathbf{R}_\sigma, \quad \prod_{\sigma'} \mathbf{B}_i \triangleleft \mathbf{R}_{\sigma'}. \quad (5.4)$$

By the Linking Lemma, $\mathbf{R} \leq_{\text{sd}} \mathbf{R}_k \times \mathbf{R}_{k'}$ is linked for some k . Thus, by Lemma 5.1, the proof is complete. \square

Theorem 5.5 can be adapted to handle cases in which there are multiple abelian factors.

Corollary 5.6. *Let $\mathbf{A}_0, \mathbf{A}_1, \dots, \mathbf{A}_{n-1}$ be finite algebras in a Taylor variety, and let $\mathbf{B}_i \triangleleft \mathbf{A}_i$ for all $0 \leq i < n$. Suppose $\alpha \subseteq \underline{n}$ and*

- \mathbf{A}_i is abelian for each $i \in \alpha$,
- \mathbf{A}_i is nonabelian and simple for each $i \in \alpha'$,
- $\mathbf{R} \leq_{\text{sd}} \mathbf{A}_0 \times \mathbf{A}_1 \times \dots \times \mathbf{A}_{n-1}$,
- $\rho_i \neq \rho_j$, for all $i, j \in \alpha'$ with $i \neq j$, and
- $R' := R \cap (B_0 \times B_1 \times \dots \times B_{n-1}) \neq \emptyset$.

Then $\mathbf{R}' = \mathbf{R}_\alpha \times \prod_{\alpha'} \mathbf{B}_i$.

Proof. If $\alpha = \emptyset$ then the product has no abelian factors and the result follows immediately from the Rectangularity Theorem. If $\alpha \neq \emptyset$, let $\alpha' = \{i_0, i_1, \dots, i_{m-1}\}$. Clearly, $\mathbf{R} \leq_{\text{sd}} \mathbf{R}_\alpha \times \mathbf{A}_{i_0} \times \mathbf{A}_{i_1} \times \dots \times \mathbf{A}_{i_{m-1}}$, and this product has a single abelian factor, $\mathbf{R}_\alpha \leq \prod_{\alpha} \mathbf{A}_i$. If θ denotes the kernel of the projection of R onto R_α , then for all $j < m$ we have $\theta \neq \rho_{i_j}$, since \mathbf{R}_α is abelian while \mathbf{A}_{i_j} is not. We can now apply the Rectangularity Theorem to this subdirect representation since the projection congruences are pairwise distinct. \square

We conclude this subsection with two observations that facilitate application of the results above to CSP problems.

Corollary 5.7. *Let $\mathbf{A}_0, \mathbf{A}_1, \dots, \mathbf{A}_{n-1}$ be finite algebras in a Taylor variety with $\mathbf{B}_i \triangleleft \mathbf{A}_i$ for all $0 \leq i < n$, and let $\alpha \subseteq \underline{n}$. Suppose \mathbf{R} and \mathbf{S} are subdirect products of $\prod_{\underline{n}} \mathbf{A}_i$ and*

- (1) \mathbf{A}_i is abelian for each $i \in \alpha$,
- (2) \mathbf{A}_i is nonabelian and simple for each $i \notin \alpha$,
- (3) $\rho_i^R \neq \rho_j^R$ and $\rho_i^S \neq \rho_j^S$, for all $i, j \in \alpha'$ with $i \neq j$,
- (4) R and S intersect $\prod_{\underline{n}} B_i$ nontrivially, and
- (5) there exists $\mathbf{x} \in R_\alpha \cap S_\alpha$.

Then $R \cap S \neq \emptyset$.

Proof. By Corollary 5.6, $\mathbf{R}' = \mathbf{R}_\alpha \times \prod_{\alpha'} \mathbf{B}_i$ and $\mathbf{S}' = \mathbf{S}_\alpha \times \prod_{\alpha'} \mathbf{B}_i$. Therefore, $\mathbf{x} \in R_\alpha \cap S_\alpha$ implies $\{\mathbf{x}\} \times \prod_{\alpha'} \mathbf{B}_i \subseteq R \cap S$. \square

The generalization to more than two subdirect products is trivial. Nonetheless, for future reference, we record the *general formulation of the Rectangularity Theorem* in

Corollary 5.8. *Let $\mathbf{A}_0, \mathbf{A}_1, \dots, \mathbf{A}_{n-1}$ be finite algebras in a Taylor variety with $\mathbf{B}_i \triangleleft \mathbf{A}_i$ for all $0 \leq i < n$, and let $\alpha \subseteq \underline{n}$. Suppose $\mathbf{R}_\ell \leq_{\text{sd}} \prod_{\underline{n}} \mathbf{A}_i$ ($0 \leq \ell < m$) and*

- (1) \mathbf{A}_i is abelian for each $i \in \alpha$,
- (2) \mathbf{A}_i is nonabelian and simple for each $i \notin \alpha$,
- (3) $\forall \ell \in \underline{m}, \forall i, j \in \alpha', i \neq j \implies \rho_i^\ell \neq \rho_j^\ell$ (where $\rho_i^\ell := \ker(\mathbf{R}_\ell \rightarrow \mathbf{A}_i)$),
- (4) each R_ℓ intersects $\prod B_i$ nontrivially,
- (5) there exists $\mathbf{x} \in \bigcap \text{Proj}_\alpha R_\ell$.

Then $\bigcap R_\ell \neq \emptyset$. In fact $\{\mathbf{x}\} \times \prod_{i \in \alpha'} \mathbf{B}_i \subseteq \bigcap_{\ell < m} \mathbf{R}_\ell$.

6. CSP APPLICATIONS

In this section we give a precise definition of what we mean by a “constraint satisfaction problem,” and what it means for such a problem to be “tractable.” We then give examples demonstrating how to use the algebraic tools developed above to prove tractability.

6.1. Definition of a constraint satisfaction problem. We give a definition of “constraint satisfaction problem” that is convenient for our purposes. This is not the most general definition possible, but let us postpone consideration of the scope of our setup.

Let $\mathbf{A} = \langle A, \mathcal{F} \rangle$ be a finite idempotent algebra, and let $\text{Sub}(\mathbf{A})$ and $\mathbf{S}(\mathbf{A})$ denote the set of subuniverses and subalgebras of \mathbf{A} , respectively.

Definition 6.1. Let \mathfrak{A} be a collection of algebras of the same similarity type. We define $\text{CSP}(\mathfrak{A})$ to be the following decision problem:

- An n -variable instance of $\text{CSP}(\mathfrak{A})$ is a triple $\langle \mathcal{V}, \mathcal{A}, \mathcal{C} \rangle$, where
 - \mathcal{V} is a set of n variables; often we take \mathcal{V} to be $\underline{n} = \{0, 1, \dots, n-1\}$;
 - \mathcal{A} is a list $(\mathbf{A}_0, \mathbf{A}_1, \dots, \mathbf{A}_{n-1}) \in \mathfrak{A}^n$ of algebras from \mathfrak{A} , one for each variable;
 - \mathcal{C} is a list $((\sigma_0, R_0), \dots, (\sigma_{J-1}, R_{J-1}))$ of constraints; each σ_j is a scope function with arity $\text{ar}(\sigma_j) = m_j$; each R_j is a constraint relation, which we take to be the universe of a subdirect product of the algebras in \mathcal{A} whose indices belong to $\text{im } \sigma_j$; that is,

$$\mathbf{R}_j \leq_{\text{sd}} \prod_{0 \leq i < m_j} \mathbf{A}_{\sigma_j(i)}.$$

- A solution to the instance $\langle \mathcal{V}, \mathcal{A}, \mathcal{C} \rangle$ is an assignment $f: \mathcal{V} \rightarrow \bigcup_{\underline{n}} A_i$ of values to variables that satisfies all constraint relations. More precisely, $f \in \prod_{\underline{n}} A_i$ and $f \circ \sigma_j \in R_j$ holds for all $0 \leq j < J$. We denote the set of solutions to $\langle \mathcal{V}, \mathcal{A}, \mathcal{C} \rangle$ by $\text{Sol}(\mathcal{C}, \underline{n})$.

Remark 6.2.

- (i) The scope function $\sigma_i: m_i \rightarrow \mathcal{V}$ picks the m_i variables involved in the constraint relation R_i . Thus, $((\sigma_0, R_0), \dots, (\sigma_{J-1}, R_{J-1})) \in \prod_{i < J} \mathcal{V}^{m_i}$.
- (ii) Frequently we require that arities of scope functions are bounded above by, say, $m \geq \text{ar}(\sigma_j)$ ($j < J$), which yields the *local constraint satisfaction problem* $\text{CSP}(\mathfrak{A}, m)$, with instances so restricted.
- (iii) If $(\sigma, R) \in \mathcal{C}$ is a constraint of an n -variable instance, and if $\text{Sol}((\sigma, R), \underline{n})$ denotes the set of tuples in $\prod_{\underline{n}} A_i$ that satisfy (σ, R) , then

$$\text{Sol}((\sigma, R), \underline{n}) = R^{\overleftarrow{\sigma}} := \left\{ \mathbf{x} \in \prod_{i \in \underline{n}} A_i \mid \mathbf{x} \circ \sigma \in R \right\}.$$

Thus, the instance $\langle \mathcal{V}, \mathcal{A}, \mathcal{C} \rangle$ has solution set $\text{Sol}(\mathcal{C}, \underline{n}) = \bigcap_{(\sigma, R) \in \mathcal{C}} R^{\overleftarrow{\sigma}}$.

- (iv) If \mathfrak{A} contains a single algebra, we write $\text{CSP}(\mathbf{A})$ instead of $\text{CSP}(\{\mathbf{A}\})$. It is important to note that, in our definition of an instance of $\text{CSP}(\mathfrak{A})$, a constraint relation is a subdirect product of algebras in \mathfrak{A} . Thus, constraint relations of an instance of $\text{CSP}(\mathbf{A})$ are subdirect powers of \mathbf{A} . In the literature it is conventional to allow instances of $\text{CSP}(\mathbf{A})$ with constraint relations that are (not necessarily subdirect) subpowers of \mathbf{A} . However, there is a simple “arc-consistency” algorithm (see, e.g., [BK17, Sec. 4.2.1]) that takes such an instance and either shows it has no solution or transforms it into an instance that has subdirect constraint relations and the same solution set as the original instance.

6.2. Instance size and tractability. The computational complexity of an algorithm for solving instances of a CSP is measured as a function of input size. Thus, in order to say what it means for $\text{CSP}(\mathfrak{A})$ to be “computationally tractable,” we need a definition of input size. This amounts to determining the number of bits required to completely specify an instance of the problem. In practice, an upper bound on the size is usually sufficient.

Using the notation in Definition 6.1 as a guide, we bound the size of an instance $\mathcal{J} = \langle \mathcal{V}, \mathcal{A}, \mathcal{C} \rangle$ of $\text{CSP}(\mathfrak{A})$. Let $q = \max(|A_0|, |A_1|, \dots, |A_{n-1}|)$, let r be the maximum rank of an operation symbol in the similarity type, and p the number of operation symbols. Then each member of the list \mathcal{A} requires at most $pq^r \log q$ bits to specify. Thus

$$\text{size}(\mathcal{A}) \leq npq^r \log q.$$

Similarly, each constraint scope $\sigma_j: m_j \rightarrow n$ can be encoded using $m_j \log n$ bits. Taking $m = \max(m_1, \dots, m_{J-1})$, we have

$$\text{size}(\sigma_0, \sigma_1, \dots, \sigma_{J-1}) \leq Jm \log n.$$

Finally, the constraint relation R_j requires at most $q^{m_j} \cdot m_j \cdot \log q$ bits. Thus,

$$\text{size}(R_0, R_1, \dots, R_{J-1}) \leq Jq^m \cdot m \log q.$$

Combining these encodings and using the fact that $\log q \leq q$, we deduce

$$\text{size}(\mathcal{J}) \leq npq^{r+1} + Jmq^{m+1} + Jmn. \quad (6.1)$$

In particular, for the problem $\text{CSP}(\mathfrak{A}, m)$, the maximum scope arity, m , is fixed, as is r . In this case, we can assume $J \leq n^m$. Consequently, $\text{size}(\mathcal{J}) \in O((nq)^{m+1})$ which yields a polynomial bound (in nq) for the size of the instance.

A problem is called *tractable* if there exists a deterministic algorithm that solves all instances of the problem and does so in an amount of time that is bounded above by a polynomial function of the input size. We can use Definition 6.1 above to classify the complexity of an algebra \mathbf{A} , or collection of algebras \mathfrak{A} , according to the complexity of their corresponding constraint satisfaction problems.

An algorithm A is called a *polynomial-time algorithm* for $\text{CSP}(\mathfrak{A})$ if there exist constants c and d such that, given an instance \mathcal{J} of $\text{CSP}(\mathfrak{A})$ of size $S = \text{size}(\mathcal{J})$, A halts in at most cS^d steps and outputs whether or not \mathcal{J} has a solution. In this case, we say A “solves the decision problem $\text{CSP}(\mathfrak{A})$ in polynomial time” and we call the algebras in \mathfrak{A} “jointly tractable.” We say that \mathfrak{A} is *jointly locally tractable* if, for every natural number, m , there is a polynomial-time algorithm A_m that solves $\text{CSP}(\mathfrak{A}, m)$. We call an algebra \mathbf{A} *tractable* when $\mathfrak{A} = \mathbf{S}(\mathbf{A})$ is jointly tractable.

We emphasize that $\text{CSP}(\mathfrak{A})$ is a *decision problem*, that is, the algorithm is only required to respond “yes” or “no” to the question of whether a particular instance has a solution, it does not have to actually produce a solution. However, it turns out that if $\text{CSP}(\mathfrak{A})$ is tractable then the corresponding *search problem* is also tractable; thus, one could design the algorithm to find a solution in polynomial time, if a solution exists, see [BJK05, Cor 4.9].

6.3. Sufficient conditions for tractability. A *meet semidistributive* lattice is one that satisfies the quasiidentity $x \wedge y \approx x \wedge z \rightarrow x \wedge y \approx x \wedge (y \vee z)$. A variety is called $\text{SD-}\wedge$ if every member of the variety has a meet semidistributive congruence lattice. Idempotent $\text{SD-}\wedge$ varieties are known to be Taylor varieties [HM88]. In [BK12], Barto and Kozik proved the following.

Theorem 6.3 ([BK12]). *A finite idempotent algebra in an SD- \wedge variety is tractable.*

A second important technique for establishing tractability is the “few subpowers algorithm” of [IMM⁺10], which its discoverers describe as a broad generalization of Gaussian elimination.

Definition 6.4. Let \mathcal{V} be a variety and k an integer, $k > 1$. A $(k + 1)$ -ary term t is called a k -edge term for \mathcal{V} if the following k identities hold in \mathcal{V} :

$$\begin{aligned} t(y, y, x, x, x, \dots, x) &\approx x \\ t(y, x, y, x, x, \dots, x) &\approx x \\ t(x, x, x, y, x, \dots, x) &\approx x \\ &\vdots \\ t(x, x, x, x, x, \dots, y) &\approx x. \end{aligned}$$

Clearly, every Mal’tsev term and every near unanimity term is an edge term, and every edge term is idempotent and Taylor. Combining the main results of [BIM⁺10] and [IMM⁺10] yields the following theorem.

Theorem 6.5. *Let \mathbf{A} be a finite idempotent algebra with an edge term. Then \mathbf{A} is tractable.*

Finally, we note that tractability is largely preserved by familiar algebraic constructions.

Theorem 6.6 ([BJK05]). *Let \mathbf{A} be a finite, idempotent, tractable algebra. Every subalgebra and finite power of \mathbf{A} is tractable. Every homomorphic image of \mathbf{A} is locally tractable.*

6.4. Rectangularity Theorem: obstacles and applications. The goal of this subsection is to consider aspects of the Rectangularity Theorem that limit its utility as a tool for proving tractability of CSPs. We give a brief overview of the potential obstacles, and then consider each one in more detail.

- (1) **Abelian factors must have easy partial solutions.** One potential limitation concerns abelian factors in the product algebra associated with a CSP instance. Indeed, application of Corollary 5.8 requires a partial solution $\mathbf{x} \in \bigcap \text{Proj}_\alpha R_\ell$ to the restricted instance derived from projecting constraints onto abelian factors. Fortunately, Section 6.4.1 will show that this obstacle is easily overcome.
- (2) **Intersecting products of minimal absorbing subalgebras.** The Rectangularity Theorem (and corollaries) assumes the universes of the subdirect products in question intersect nontrivially with a single product $\prod B_i$ of minimal absorbing subuniverses. We refer to “minimal absorbing subuniverses” quite frequently, so from now on we call them *masses*; that is,
 - *mass* := minimal absorbing subuniverse;
 - *mass product* := product of masses.
 Assumption (4) of Corollary 5.8 requires all constraint relations intersect nontrivially with a single mass product. This is a real limitation, as we demonstrate in Section 6.4.2.
- (3) **Nonabelian factors must be simple.** This is the most obvious limitation of the theorem and at this point we don’t have a completely general means of overcoming it. However, there are some work-arounds that can be useful in special cases; we describe these below.

In the next two subsections we address potential limitations (1) and (2). In Section 7 we develop alternative methods for proving tractability of nonsimple algebras, and in Section 8 we apply these methods in the special setting of “commutative idempotent binars.”

6.4.1. *Tractability of abelian algebras.* To address concern (1) above we observe that finite abelian algebras yield tractable CSPs. We show how to use this and the Rectangularity Theorem to find solutions to CSP instances (or prove none exists). To begin, recall the result of tame congruence theory [HM88, Thm 7.12] that we reformulated above in Lemma 3.4. As noted in Remark 3.6 above, this result has the following important corollary. (For a proof that avoids tame congruence theory, see [BKS15, Thm 5.1].)

Theorem 6.7. *Let \mathcal{V} be a locally finite variety with a Taylor term. Every finite abelian member of \mathcal{V} generates a congruence-permutable variety. Consequently, every finite abelian member of \mathcal{V} is tractable.*

Let \mathbf{A} be a finite algebra in a Taylor variety and fix an instance $\mathcal{J} = \langle \mathcal{V}, \mathcal{A}, \mathcal{C} \rangle$ of $\text{CSP}(\mathcal{S}(\mathbf{A}))$ with $n = |\mathcal{V}|$ variables. Suppose all nonabelian algebras in the list \mathcal{A} are simple. Let $\alpha \subseteq \underline{n}$ denote the indices of the abelian algebras in \mathcal{A} , and assume without loss of generality that $\alpha = \{0, 1, \dots, q-1\}$. That is, $\mathbf{A}_0, \mathbf{A}_1, \dots, \mathbf{A}_{q-1}$ are finite idempotent abelian algebras. Consider the restricted instance \mathcal{J}_α obtained by dropping all constraints with scopes that do not intersect α , and by restricting the remaining constraint relations to the abelian factors. Since the only algebras involved in \mathcal{J}_α are abelian, this is an instance of a tractable CSP. Therefore, we can check in polynomial-time whether or not \mathcal{J}_α has a solution. If there is no solution, then the original instance \mathcal{J} has no solution. On the other hand, suppose $f_\alpha \in \prod_{j \in \alpha} A_j$ is a solution to \mathcal{J}_α . In Corollary 5.8, to reach the conclusion that the full instance has a solution, we required a partial solution $\mathbf{x} \in \bigcap \text{Proj}_\alpha R_\ell$. This is precisely what f_α provides.

To summarize, we look for a partial solution by restricting the instance to abelian factors and, if successful, we use the partial solution for \mathbf{x} in Corollary 5.8. Then, if the remaining hypotheses of Corollary 5.8 hold, we conclude that a solution to the original instance exists. If no solution to the restricted instance exists, then the original instance has no solution. Thus, assumption (5) of Corollary 5.8 does not limit the application of the result.

6.4.2. *Mass products.* This section concerns products of minimal absorbing subalgebras, or “mass products.” Hypothesis (4) of Corollary 5.8 assumes that all constraint relations intersect nontrivially with a single mass product. Instances where this hypothesis does not hold are easy to contrive. For example, take the algebra $\mathbf{A} = \langle \{0, 1\}, m \rangle$, where $m: A^3 \rightarrow A$ is an idempotent majority operation—that is, $m(x, x, x) \approx x$ and $m(x, x, y) \approx m(x, y, x) \approx m(y, x, x) \approx x$. Let $\mathbf{R} = \langle R, m \rangle$ and $\mathbf{S} = \langle S, m \rangle$ be subdirect products of \mathbf{A}^3 with universes $R = \{(0, 0, 0), (0, 0, 1), (0, 1, 0), (1, 0, 0)\}$ and $S = \{(0, 1, 1), (1, 0, 1), (1, 1, 0), (1, 1, 1)\}$. Then there are mass products that intersect nontrivially with either R or S , and each of R and S fully contains every mass product that it intersects. However, there is no single mass product intersecting nontrivially with both R and S .

The example described above is very special. In particular, there is no solution to the instance with constraint relations R and S . We might hope that when an instance *does* have a solution, then there must be a solution that passes through a mass product. As we

now demonstrate, this is false. In fact, Example 6.9 presents subdirect powers that intersect nontrivially (so the instance has a solution), yet their intersection avoids every mass product (so the Rectangularity Theorem cannot locate a solution nor prove none exists).

Proposition 6.8. *There exists an algebra \mathbf{A} with subdirect powers \mathbf{R} and \mathbf{S} such that $R \cap S \neq \emptyset$ and $R \cap S \cap \prod B_i = \emptyset$, for every collection $\{B_i\}$ of masses.*

We prove this by simply producing an example that meets the stated conditions.

Example 6.9. Let $\mathbf{A} = \langle \{0, 1, 2\}, \circ \rangle$ be an algebra with binary operation \circ given by

\circ	0	1	2
0	0	1	2
1	1	1	0
2	2	0	2

The proper nonempty subuniverses of \mathbf{A} are $\{0\}$, $B_1 := \{1\}$, $B_2 := \{2\}$, $\{0, 1\}$, and $\{0, 2\}$. Note that B_1 and B_2 are masses with respect to the term $t(x, y, z, w) = (x \circ y) \circ (z \circ w)$. Note also that $\{0\}$ is not an absorbing subuniverse of \mathbf{A} . For if $\{0\} \triangleleft \mathbf{A}$, then $\{0\} \triangleleft \{0, 1\}$ which is false, since $\{0, 1\}$ is a semilattice with absorbing element 1.

Let $\mathbf{A}_0 \cong \mathbf{A} \cong \mathbf{A}_1$, $R = \{(0, 0), (1, 1), (2, 2)\}$, and $S = \{(0, 0), (1, 2), (2, 1)\}$. Then \mathbf{R} and \mathbf{S} are subdirect products of $\mathbf{A}_0 \times \mathbf{A}_1$ and $R \cap S = \{(0, 0)\}$. There are four masses of $\mathbf{A}_0 \times \mathbf{A}_1$, namely, $B_i \times B_j = \{(i, j)\}$, for $i, j \in \{1, 2\}$. Finally, observe,

$$R \cap (B_i \times B_j) = \begin{cases} \{(i, j)\}, & i = j, \\ \emptyset, & i \neq j; \end{cases} \quad S \cap (B_i \times B_j) = \begin{cases} \emptyset, & i = j \\ \{(i, j)\}, & i \neq j. \end{cases}$$

Thus, $R \cap S \cap (B_i \times B_j) = \emptyset$ for all $i, j \in \{1, 2\}$, proving Proposition 6.8.

6.5. Algorithm synthesis for heterogeneous problems. We now present a new result (Theorem 6.13) which, like the Rectangularity Theorem, aims to describe some of the elements that must belong to certain subdirect products. The conclusions we draw are weaker than those of Theorem 5.5, but the hypotheses required here are simpler, and the motivation is different.

We have in mind subdirect products of “heterogeneous” families of algebras. What we mean by this is described and motivated as follows: let $\mathcal{C}_1, \dots, \mathcal{C}_m$ be classes of algebras, all of the same signature. Perhaps we are lucky enough to possess a single algorithm that proves the algebras in $\bigcup \mathcal{C}_i$ are jointly tractable. Suppose instead that we have no such single algorithm at hand, but at least we do happen to know that the classes \mathcal{C}_i are *separately tractable*; that is, for each i we have a proof (or program) P_i certifying the tractability of $\text{CSP}(\mathcal{C}_i)$. It is natural to ask whether and under what conditions it might be possible to derive from $\{P_i \mid 1 \leq i \leq m\}$ a single proof (program) certifying that the algebras in $\bigcup \mathcal{C}_i$ are *jointly tractable*, so that $\text{CSP}(\bigcup \mathcal{C}_i)$ is tractable.

The results in this section take a step in this direction by considering two special classes of algebras that are known to be separately tractable, and demonstrating that they are, in fact, jointly tractable. We apply this tool in Section 7.2.2 to prove tractability of a CSP involving algebras that were already known to be tractable, but not previously known to be jointly tractable.

The results here involve special term operations called cube terms and transitive terms. A k -ary idempotent term t is a *cube term* if for every coordinate $i \leq k$, t satisfies an identity of the form $t(x_1, \dots, x_k) \approx y$, where $x_1, \dots, x_k \in \{x, y\}$ and $x_i = x$. (See [BIM⁺10].) A k -ary operation f on a set A is called *transitive in the i -th coordinate* if for every $u, v \in A$, there exist $a_1, \dots, a_k \in A$ such that $a_i = u$ and $f(a_1, \dots, a_k) = v$. Operations that are transitive in every coordinate are called *transitive*. (See [BKS15].)

Fact 6.10. Let \mathbf{A} be a finite idempotent algebra and suppose t is a cube term operation on \mathbf{A} . Then t is a transitive term operation on \mathbf{A} .

Proof. This follows immediately from the definitions. \square

Fact 6.11. The class

$$\mathcal{T} = \{\mathbf{A} \mid \mathbf{A} \text{ finite and every subalgebra of } \mathbf{A} \text{ has a transitive term op}\} \quad (6.2)$$

is closed under the taking of homomorphic images, subalgebras, and finite products. That is, \mathcal{T} is a *pseudovariety*.

We also require the following obvious fact about nontrivial algebras in a Taylor variety. (We call an algebra *nontrivial* if it has more than one element in its universe.)

Fact 6.12. If \mathbf{A} and \mathbf{B} are nontrivial algebras in a Taylor variety \mathcal{V} , then for some $k > 1$ there is a k -ary term t in \mathcal{V} such that $t^{\mathbf{A}}$ and $t^{\mathbf{B}}$ each depends on at least two of its arguments.

Finally, we are ready for the main result of this section.

Theorem 6.13. Let $\mathbf{A}_0, \mathbf{A}_1, \dots, \mathbf{A}_{n-1}$ be finite idempotent algebras in a Taylor variety and assume there exists a proper nonempty subset $\alpha \subset \underline{n}$ such that

- \mathbf{A}_i has a cube term for all $i \in \alpha$,
- \mathbf{A}_j has a sink $s_j \in A_j$ for all $j \in \alpha'$; let $\mathbf{s} \in \prod_{\alpha'} A_j$ be a tuple of sinks.

If $\mathbf{R} \leq_{\text{sd}} \prod_{\underline{n}} \mathbf{A}_i$, then the set $X := R_\alpha \times \{\mathbf{s}\}$ is a subuniverse of \mathbf{R} .

Remark 6.14. To foreshadow applications of Theorem 6.13, imagine we have algebras of the type described, and a collection \mathcal{R} of subdirect products of these algebras. Suppose also that we have somehow determined that the intersection of the α -projections of these subdirect products is nonempty, say, $\mathbf{x}_\alpha \in \bigcap_{R \in \mathcal{R}} R_\alpha$. Then the full intersection $\bigcap \mathcal{R}$ will also be nonempty, since according to the theorem it must contain the tuple that is \mathbf{x}_α on α and \mathbf{s} off α .

Proof. Fix $\mathbf{x} \in X$, so $\mathbf{x} \circ \alpha' = \mathbf{s}$ and $\mathbf{x} \circ \alpha \in R_\alpha$. We prove $\mathbf{x} \in R$. Define $\mathbf{A}_\alpha = \prod_{\alpha} \mathbf{A}_i$ and $\mathbf{A}_{\alpha'} = \prod_{\alpha'} \mathbf{A}_i$, so R_α and R are subuniverses of \mathbf{A}_α and $\mathbf{A}_\alpha \times \mathbf{A}_{\alpha'}$, respectively.

By Fact 6.10, for each $i \in \alpha$ every subalgebra of \mathbf{A}_i has a transitive term operation. Moreover, the class \mathcal{T} defined in (6.2) is a pseudovariety, so every subalgebra of \mathbf{A}_α has a transitive term operation. Suppose there are J subalgebras of \mathbf{A}_α and let $\{t_j \mid 0 \leq j < J\}$ denote the collection of transitive term operations, one for each subalgebra. Then it is not hard to prove that $t := t_0 \star t_1 \star \dots \star t_{J-1}$ (defined in Sec. 4.1) is a transitive term for every subalgebra of \mathbf{A}_α . In particular t is transitive for the subalgebra with universe R_α .

Assume t is N -ary. Fix $j \in \alpha'$. We may assume without loss of generality that $t^{\mathbf{A}_j}$ depends on its first argument, at position 0. (It must depend on at least one argument by idempotence.) Now, since \mathbf{R} is a subdirect product of $\prod_{\underline{n}} \mathbf{A}_i$, there exists $\mathbf{r}^{(j)} \in R$ such

that $\mathbf{r}^{(j)}(j) = s_j$, the sink in \mathbf{A}_j . Since $\mathbf{x} \circ \alpha \in R_\alpha$ and since $t^{\mathbf{A}_\alpha}$ is transitive over \mathbf{R}_α , there exist $\mathbf{r}_1, \dots, \mathbf{r}_{N-1}$ in R such that

$$\begin{aligned} \mathbf{y}^{(j)} &:= t^{\mathbf{A}_\alpha \times \mathbf{A}_{\alpha'}}(\mathbf{r}^{(j)}, \mathbf{r}_1, \dots, \mathbf{r}_{N-1}) \\ &= (t^{\mathbf{A}_\alpha}(\mathbf{r}^{(j)} \circ \alpha, \mathbf{r}_1 \circ \alpha, \dots, \mathbf{r}_{N-1} \circ \alpha), t^{\mathbf{A}_{\alpha'}}(\mathbf{r}^{(j)} \circ \alpha', \mathbf{r}_1 \circ \alpha', \dots, \mathbf{r}_{N-1} \circ \alpha')) \end{aligned}$$

belongs to R and satisfies $\mathbf{y}^{(j)} \circ \alpha = \mathbf{x} \circ \alpha$ and $\mathbf{y}^{(j)}(j) = s_j$.

If $\alpha' = \{j\}$, we are done, since $\mathbf{y}^{(j)} = \mathbf{x}$ in that case. If $|\alpha'| > 1$, then we repeat the foregoing procedure for each $j \in \alpha'$ and obtain a subset $\{\mathbf{y}^{(j)} \mid j \in \alpha'\}$ of R , each member of which agrees with \mathbf{x} on α and has a sink in some position $j \in \alpha'$.

Next, choose distinct $j, k \in \alpha'$. Suppose w is a Taylor term for \mathcal{V} . Then by Fact 6.12 we may assume without loss of generality that $w^{\mathbf{A}_j}$ depends on its p -th argument and $w^{\mathbf{A}_k}$ depends on its q -th argument, for some $p \neq q$. Consider

$$\mathbf{z} := w^{\Pi \mathbf{A}_i}(\mathbf{y}^{(j)}, \dots, \mathbf{y}^{(j)}, \mathbf{y}^{(k)}, \mathbf{y}^{(j)}, \dots, \mathbf{y}^{(j)})$$

$\widehat{}$
 \widehat{q} -th argument

Evidently, $\mathbf{z}(j) = s_j$, $\mathbf{z}(k) = s_k$, and $\mathbf{z} \circ \alpha = \mathbf{x} \circ \alpha$ by idempotence, since, when restricted to indices in α , all the input arguments agree and are equal to $\mathbf{x} \circ \alpha$. If $\alpha' = \{j, k\}$, we are done. Otherwise, choose $\ell \in \alpha' - \{j, k\}$, and again $w^{\mathbf{A}_\ell}$ depends on at least one of its arguments, say, the u -th. Let

$$\mathbf{z}' := w^{\Pi \mathbf{A}_i}(\mathbf{z}, \dots, \mathbf{z}, \mathbf{y}^{(\ell)}, \mathbf{z}, \dots, \mathbf{z}).$$

$\widehat{}$
 \widehat{u} -th argument

Then \mathbf{z}' belongs to R , agrees with \mathbf{x} on α , and satisfies $\mathbf{z}'(j) = s_j$, $\mathbf{z}'(k) = s_k$, and $\mathbf{z}'(\ell) = s_\ell$. Continuing in this way until the set α' is exhausted produces an element in R that agrees with \mathbf{x} everywhere. In other words, \mathbf{x} itself belongs to R . \square

In Section 7.2.2 we apply Theorem 6.13 in the special case where “has a cube term” in the first hypothesis is replaced with “is abelian.” Let us be explicit.

Corollary 6.15. *Let $\mathbf{A}_1, \dots, \mathbf{A}_{n-1}$ be finite idempotent algebras in a locally finite Taylor variety. Suppose there exists $0 < k < n - 1$ such that*

- \mathbf{A}_i is abelian for all $i < k$;
- \mathbf{A}_i has a sink $s_i \in A_i$ for all $i \geq k$.

If $\mathbf{R} \leq_{\text{sd}} \prod \mathbf{A}_i$, then $R_k \times \{s_k\} \times \{s_{k+1}\} \times \dots \times \{s_{n-1}\} \subseteq R$.

Proof. Since $\mathbf{A}_k := \prod_{i < k} \mathbf{A}_i$ is abelian and lives in a locally finite Taylor variety, there exists a term m such that $m^{\mathbf{A}_k}$ is a Mal'tsev operation on \mathbf{A}_k (Theorem 6.7). Since a Mal'tsev term is a cube term, the result follows from Theorem 6.13. \square

7. PROBLEM INSTANCE REDUCTIONS

In this section we develop some useful notation for taking an instance of a CSP and restricting or reducing it in various ways, either by removing variables or by reducing modulo a sequence of congruence relations. The utility of these tools will be demonstrated in Section 8.

Throughout this section, \mathbf{A} will denote a finite idempotent algebra. The problem we will focus on is $\text{CSP}(\mathfrak{A})$, defined in Section 6.1, and we will be particularly interested in the special case in which $\mathfrak{A} = \mathbf{S}(\mathbf{A})$.

Recall, we denote an n -variable instance of $\text{CSP}(\mathbf{S}(\mathbf{A}))$ by $\mathcal{J} = \langle \underline{n}, \mathcal{A}, \mathcal{C} \rangle$, where $\underline{n} = \{0, 1, \dots, n-1\}$ represents the set of variables, $\mathcal{A} = (\mathbf{A}_0, \mathbf{A}_1, \dots, \mathbf{A}_{n-1})$ is a list of n subalgebras of \mathbf{A} , and $\mathcal{C} = ((\sigma_0, R_0), (\sigma_1, R_1), \dots, (\sigma_{J-1}, R_{J-1}))$ is a list of constraints with respective arities $\text{ar}(\sigma_j) = m_j$. Thus, $R_j \subseteq \prod_{i < m_j} A_{\sigma(i)}$. Much of the discussion below refers to an arbitrary constraint in \mathcal{C} . In such cases it will simplify notation to drop subscripts and denote the constraint by (σ, R) .

7.1. Variable reductions.

7.1.1. Partial scopes and partial constraints. Consider the restriction of an n -variable CSP instance \mathcal{J} to the first k of its variables, for some $k \leq n$. This results in a new *partial scope* given by the function $\sigma|_{\sigma^{-1}(\underline{k} \cap \text{im } \sigma)}$. Call this the *k -partial scope of σ* and, to simplify the notation, let

$$\sigma|_{\underline{k}} = \sigma|_{\sigma^{-1}(\underline{k} \cap \text{im } \sigma)}.$$

If $\underline{k} \cap \text{im } \sigma = \emptyset$ then the k -partial scope of σ is the empty function. To obtain the *k -partial constraint of (σ, R)* , we take the k -partial scope of σ as the new scope; for the constraint relation we take the restriction of each tuple in R to its first $p = |\underline{k} \cap \text{im } \sigma|$ coordinates. If we let $R|_{\underline{k}} = R|_{\sigma^{-1}(\underline{k} \cap \text{im } \sigma)}$, then the *k -partial constraint of (σ, R)* is given by $(\sigma|_{\underline{k}}, R|_{\underline{k}})$.

For example, let σ be a scope consisting of variables 2, 4, 7, so $(\sigma(0), \sigma(1), \sigma(2)) = (2, 4, 7)$. To find the 5-partial constraint of (σ, R) , restrict (σ, R) to the first $k = 5$ variables of the instance. Then, $\sigma^{-1}(\underline{k} \cap \text{im } \sigma) = \sigma^{-1}\{2, 4\} = \{0, 1\}$, and $\sigma|_{\underline{k}} = (\sigma(0), \sigma(1)) = (2, 4)$, so $R|_{\underline{k}} = \{(x, y) \mid (x, y, z) \in R\}$.

7.1.2. Partial instances. The *k -partial instance $\mathcal{J}_{\underline{k}}$* of \mathcal{J} is the restriction of \mathcal{J} to its first k variables. Thus, $\mathcal{J}_{\underline{k}}$ is the instance with constraint set $\mathcal{C}_{\underline{k}}$ equal to the set of all k -partial constraints of \mathcal{J} . If we let $\text{Sol}(\mathcal{J}, \underline{k})$ denote the set of solutions to $\mathcal{J}_{\underline{k}}$, then $f \in \text{Sol}(\mathcal{J}, \underline{k})$ if and only if for each $j \in J$, we have $f \circ \sigma_j|_{\underline{k}} \in R_j|_{\underline{k}}$. We might be tempted to call $\text{Sol}(\mathcal{J}, \underline{k})$ the set of “partial solutions,” but that’s a bit misleading since an $f \in \text{Sol}(\mathcal{J}, \underline{k})$ may or may not extend to a solution to the full instance \mathcal{J} .

7.2. Quotient reductions. Let \mathbf{A} be a finite idempotent algebra and $\mathcal{J} = \langle \underline{n}, \mathcal{A}, \mathcal{C} \rangle$ an n -variable instance of $\text{CSP}(\mathbf{S}(\mathbf{A}))$. As above, we assume $\mathcal{A} = (\mathbf{A}_0, \mathbf{A}_1, \dots, \mathbf{A}_{n-1}) \in \mathbf{S}(\mathbf{A})^n$, $\mathcal{C} = ((\sigma_0, R_0), (\sigma_1, R_1), \dots, (\sigma_{J-1}, R_{J-1}))$, and $\mathbf{R}_j \leq_{\text{sd}} \prod_{m_j} A_{\sigma(i)}$, for $0 \leq j < J$.

7.2.1. *Quotient instances.* Suppose $\theta_i \in \text{Con}(\mathbf{A}_i)$ ($0 \leq i < n$) and define $\boldsymbol{\theta} = (\theta_0, \dots, \theta_{n-1})$. For $\mathbf{x} = (x_0, x_1, \dots, x_{n-1}) \in \prod_n A_i$, let $\mathbf{x}/\boldsymbol{\theta}$ denote the tuple whose i -th component is the θ_i -class of \mathbf{A}_i containing x_i , so $\mathbf{x}/\boldsymbol{\theta} = (x_0/\theta_0, x_1/\theta_1, \dots, x_{n-1}/\theta_{n-1}) \in \prod_{i \in \underline{n}} A_i/\theta_i$.

We need a more general notation than this since our tuples may come from $\prod_m A_{\sigma(i)}$ for some scope function $\sigma: m \rightarrow n$. Viewing $\mathbf{x} \in \prod_n A_i$ (resp., $\boldsymbol{\theta}$) as a function from n to $\bigcup A_i$ (resp., $\bigcup \text{Con}(\mathbf{A}_i)$), write $\mathbf{x} \circ \sigma \in \prod_m A_{\sigma(i)}$ and $\boldsymbol{\theta} \circ \sigma \in \prod_m \text{Con}(\mathbf{A}_{\sigma(i)})$ for the corresponding scope-restricted tuples, and then define

$$(\mathbf{x} \circ \sigma)/(\boldsymbol{\theta} \circ \sigma) = (x_{\sigma(0)}/\theta_{\sigma(0)}, x_{\sigma(1)}/\theta_{\sigma(1)}, \dots, x_{\sigma(m-1)}/\theta_{\sigma(m-1)}) \in \prod_{i \in \underline{m}} A_{\sigma(i)}/\theta_{\sigma(i)}.$$

Given $\boldsymbol{\theta}$ and $(\sigma, R) \in \mathcal{C}$, define the set $\mathcal{C}/\boldsymbol{\theta}$ of *quotient constraints* by $\{(\sigma, R/\boldsymbol{\theta}) \mid (\sigma, R) \in \mathcal{C}\}$, where

$$R/\boldsymbol{\theta} := \{\mathbf{r}/(\boldsymbol{\theta} \circ \sigma) \in \prod_{i \in \underline{m}} A_{\sigma(i)}/\theta_{\sigma(i)} \mid \mathbf{r} \in R\}. \quad (7.1)$$

Here are a few easily proved facts about quotient constraints.

Fact 7.1. If $\boldsymbol{\theta} \in \prod_n \text{Con}(\mathbf{A}_i)$ and if $\mathbf{R} \leq_{\text{sd}} \prod_m \mathbf{A}_{\sigma(i)}$, then $R/\boldsymbol{\theta}$ defined in (7.1) is a subuniverse of $\prod_m \mathbf{A}_{\sigma(i)}/\theta_{\sigma(i)}$ and the corresponding subalgebra is subdirect.

Fact 7.2. If $\boldsymbol{\theta} \in \prod_n \text{Con}(\mathbf{A}_i)$ and if \mathcal{J} is an n -variable instance of $\text{CSP}(\mathbf{A})$, then the constraint set $\mathcal{C}/\boldsymbol{\theta}$ defines an n -variable instance of $\text{CSP}(\text{HS}(\mathbf{A}))$.

The *quotient instance* $\mathcal{J}/\boldsymbol{\theta}$ is the instance of $\text{CSP}(\text{HS}(\mathbf{A}))$ described in Fact 7.2.

Fact 7.3. If \mathbf{x} is a solution to \mathcal{J} , then $\mathbf{x}/\boldsymbol{\theta}$ is a solution to $\mathcal{J}/\boldsymbol{\theta}$.

By Fact 7.3, if there is a quotient instance with no solution, then \mathcal{J} has no solution. However, there may be solutions to every proper quotient instance and yet no solution to \mathcal{J} .

7.2.2. *Block instances.* Let \mathfrak{A} be a collection of finite idempotent algebras. Let $\mathcal{J} = \langle \underline{n}, \mathcal{A}, \mathcal{C} \rangle$ be an n -variable instance of $\text{CSP}(\mathfrak{A})$, where $\mathcal{A} = (\mathbf{A}_0, \mathbf{A}_1, \dots, \mathbf{A}_{n-1})$ and \mathcal{C} is a finite set of constraints. If $\boldsymbol{\theta} \in \prod_n \text{Con}(\mathbf{A}_i)$ and $\mathbf{x} \in \prod_n A_i$, then the tuple $\mathbf{x}/\boldsymbol{\theta}$ of θ_i -blocks is actually a list of algebras, by idempotence.

For each constraint $(\sigma, R) \in \mathcal{C}$, consider restricting the relation R to the x_i/θ_i -classes in its scope σ . In other words, replace the constraint (σ, R) with the *block constraint* $(\sigma, R \cap \Pi_{\sigma} \mathbf{x}/\boldsymbol{\theta})$, where we have defined $\Pi_{\sigma} \mathbf{x}/\boldsymbol{\theta}$ to be the product of the blocks $x_{\sigma(i)}/\theta_{\sigma(i)}$ ($1 \leq i \leq m$). Finally, let $\mathcal{J}_{\mathbf{x}/\boldsymbol{\theta}} = \langle \underline{n}, \mathbf{x}/\boldsymbol{\theta}, \mathcal{C}_{\mathbf{x}/\boldsymbol{\theta}} \rangle$ denote the problem instance of $\text{CSP}(\mathcal{S}(\mathcal{A}))$ specified by the constraint set $\mathcal{C}_{\mathbf{x}/\boldsymbol{\theta}} := \{(\sigma, R \cap \Pi_{\sigma} \mathbf{x}/\boldsymbol{\theta}) \mid (\sigma, R) \in \mathcal{C}\}$. We call $\mathcal{J}_{\mathbf{x}/\boldsymbol{\theta}}$ the *$\mathbf{x}/\boldsymbol{\theta}$ -block instance* of \mathcal{J} . Evidently, a solution to $\mathcal{J}_{\mathbf{x}/\boldsymbol{\theta}}$ is also a solution to \mathcal{J} .

7.3. The Quotient-Block strategy. The notions “quotient instance” and “block instance” suggest the following strategy for solving CSPs that works in certain special cases (one of which we will see in Example 7.4): First search for a solution $\mathbf{x}/\boldsymbol{\theta}$ to the quotient instance $\mathcal{J}/\boldsymbol{\theta}$ for some conveniently chosen $\boldsymbol{\theta}$. If no quotient solution exists, then the original instance \mathcal{J} has no solution. Otherwise, if $\mathbf{x}/\boldsymbol{\theta}$ is a solution to $\mathcal{J}/\boldsymbol{\theta}$, then try to solve the $\mathbf{x}/\boldsymbol{\theta}$ -block instance of \mathcal{J} . If successful, then the instance \mathcal{J} has a solution. Otherwise, try again with a different solution to $\mathcal{J}/\boldsymbol{\theta}$. If all solutions to the quotient instance are exhausted without finding a solution to a corresponding block instance, then \mathcal{J} has no solution.

·	0	1	2	3
0	0	0	3	2
1	0	1	3	2
2	3	3	2	1
3	2	2	1	3

· Sq ₃	0	1	2
0	0	2	1
1	2	1	0
2	1	0	2

<i>t</i>	0	1	2	3
0	0	0	0	0
1	0	1	1	1
2	2	2	2	2
3	3	3	3	3

Figure 1: Operation tables of Example 7.4: for the algebra \mathbf{A} (left); for the quotient $\mathbf{A}/\Theta \cong \mathbf{Sq}_3$ (middle); for $t(x, y) = x \cdot (y \cdot (x \cdot y))$ (right).

This approach is effective as long as we can find a quotient instance \mathcal{J}/θ for which there is a polynomial bound on the number of quotient solutions. Although this is not always possible (consider instances involving only simple algebras!), there are situations in which an appropriate choice of congruences makes it easy to check that every instance has quotient instance with a very small number of solutions. We present an example.

Example 7.4. Let $\mathbf{A} = \langle \{0, 1, 2, 3\}, \cdot \rangle$, be an algebra with a single binary operation given by the table on the left in Figure 1. The proper nonempty subuniverses are $\{0, 1\}$, $\{1, 2, 3\}$, and the singletons. The algebra \mathbf{A} has a single proper nontrivial congruence relation, Θ , with partition $|01|2|3|$. The quotient algebra \mathbf{A}/Θ is a 3-element Steiner quasigroup, which happens to be abelian. Denote the latter by \mathbf{Sq}_3 and note the middle table in Figure 1 describing its binary operation. Observe that the subalgebra $\{1, 2, 3\}$ of \mathbf{A} is also isomorphic to \mathbf{Sq}_3 .

As Theorem 6.7 predicts, the algebra \mathbf{A}/Θ has a Mal'tsev term $q(x, y, z) = y \cdot (x \cdot z)$. Let $s(x, y) = q(x, y, y) = y \cdot (x \cdot y)$. Then $\mathbf{A}/\Theta \models s(x, y) \approx x$. By iterating the term s on its second variable, we arrive at a term $t(x, y)$ such that, in \mathbf{A} , $t(x, t(x, y)) = t(x, y)$. In fact, $t(x, y) = x \cdot (y \cdot (x \cdot y))$. To summarize

$$\begin{aligned} \mathbf{A} &\models t(x, t(x, y)) \approx t(x, y) \\ \mathbf{A}/\Theta &\models t(x, y) \approx x. \end{aligned}$$

The table for t appears on the right in Figure 1.

For the next result, we denote the two-element semilattice by $\mathbf{S}_2 = \langle \{0, 1\}, \wedge \rangle$.

Lemma 7.5. *Let $\mathfrak{A} = \{\mathbf{S}_2, \mathbf{Sq}_3\}$. Then $\text{CSP}(\mathfrak{A})$ is tractable.*

Proof. Let $\mathcal{J} = \langle \underline{n}, \mathcal{A}, \mathcal{C} \rangle$ be an instance of $\text{CSP}(\mathfrak{A})$. Recall, the set of solutions to \mathcal{J} is $\text{Sol}(\mathcal{C}, \underline{n}) = \bigcap_{\mathcal{C}} R^{\overline{\sigma}}$. We shall apply Corollary 6.15 to establish tractability. We have $\mathcal{A} = (\mathbf{A}_0, \mathbf{A}_1, \dots, \mathbf{A}_{n-1}) \in \mathfrak{A}^n$. Let $\alpha = \{i : \mathbf{A}_i \cong \mathbf{Sq}_3\}$ and $\alpha' = \{i : \mathbf{A}_i \cong \mathbf{S}_2\}$, so $\underline{n} = \alpha \cup \alpha'$. We may assume without loss of generality that, for each constraint (σ, R) appearing in \mathcal{C} , the associated algebra \mathbf{R} is a subdirect product of $\prod_m \mathbf{A}_{\sigma(i)}$, where $\mathbf{A}_{\sigma(i)} \cong \mathbf{Sq}_3$ for all $\sigma(i) \in \alpha$, and $\mathbf{A}_{\sigma(i)} \cong \mathbf{S}_2$ for all $\sigma(i) \in \alpha'$.

Let 0 denote the bottom element of each semilattice. For each constraint (σ, R) , let $R|_{\overline{\alpha}}$ be an abbreviation for $R|_{\sigma^{-1}(\alpha \cap \text{im } \sigma)}$ (the projection of R onto factors in its scope with indices in α). Let $\mathbf{0}$ denote a tuple of 0's of length $|\alpha' \cap \text{im } \sigma|$. Then Corollary 6.15 implies that for each constraint $(\sigma, R) \in \mathcal{C}$ we have

$$R|_{\overline{\alpha}} \times \{\mathbf{0}\} \subseteq R. \quad (7.2)$$

Obviously, if \mathcal{J} has a solution, then the α -partial instance \mathcal{J}_α also has a solution. Conversely, if $f \in \prod_\alpha A_i$ is a solution to \mathcal{J}_α , then (7.2) implies that $g \in \prod_{\underline{n}} A_i$ defined by

$$g(i) = \begin{cases} f(i), & \text{if } i \in \alpha, \\ 0, & \text{if } i \in \alpha', \end{cases}$$

is a solution to \mathcal{J} . We conclude that \mathcal{J} has a solution if and only if the α -partial instance has a solution. Since abelian algebras yield tractable CSPs, the proof is complete. \square

Now we address the tractability of the example at hand.

Proposition 7.6. *If $\langle \mathbf{A}, \cdot \rangle$ is the four-element algebra with operation table given in Figure 1, then $\text{CSP}(\mathbf{S}(\mathbf{A}))$ is tractable.*

Proof. Let $\mathcal{J} = \langle \underline{n}, \mathcal{A}, \mathcal{C} \rangle$ be an instance of $\text{CSP}(\mathbf{S}(\mathbf{A}))$ where $\mathcal{A} = (\mathbf{A}_0, \mathbf{A}_1, \dots, \mathbf{A}_{n-1}) \in \mathbf{S}(\mathbf{A})^n$. For each $i \in \underline{n}$, let $\theta_i = \Theta$ if $A_i = A$, and let $\theta_i = 0_{A_i}$ otherwise. The universes A_i/θ_i of the quotient algebras satisfy

$$A_i/\theta_i = \begin{cases} \{\{0, 1\}, \{2\}, \{3\}\}, & \text{if } A_i = \{0, 1, 2, 3\}, \\ \{\{1\}, \{2\}, \{3\}\}, & \text{if } A_i = \{1, 2, 3\}, \\ \{\{0\}, \{1\}\}, & \text{if } A_i = \{0, 1\}. \end{cases}$$

In the first two cases \mathbf{A}_i/θ_i is isomorphic to \mathbf{Sq}_3 and in the third case \mathbf{A}_i/θ_i is a 2-element meet semilattice. (Singleton factors, $\mathbf{A}_i/\theta_i = \{a\}$, are safely ignored because all solutions must assign the the value a to the associated variable.)

Let $\boldsymbol{\theta} = (\theta_0, \theta_1, \dots, \theta_{n-1})$ and consider the quotient instance $\mathcal{J}/\boldsymbol{\theta}$ which, by the observations in the preceding paragraph, is an instance of $\text{CSP}(\mathfrak{A})$, where $\mathfrak{A} = \{\mathbf{Sq}_3, \mathbf{S}_2\}$. By Lemma 7.5, $\mathcal{J}/\boldsymbol{\theta}$ can be solved in polynomial-time. If $\mathcal{J}/\boldsymbol{\theta}$ has no solution, then neither does \mathcal{J} . Otherwise, let $f \in \text{Sol}(\mathcal{J}/\boldsymbol{\theta}, \underline{n})$ be a solution to the quotient instance. We will use f to construct a solution $g \in \text{Sol}(\mathcal{J}, \underline{n})$ to the original instance. (In so doing, we often map a singleton set to its inhabitant via the usual set union operation; e.g., $\bigcup\{x\} = x$.)

Assume $\underline{n} = \alpha \cup \beta \cup \gamma$ is a disjoint union where

$$A_i = \begin{cases} \{0, 1, 2, 3\}, & \text{if } i \in \alpha, \\ \{1, 2, 3\}, & \text{if } i \in \beta, \\ \{0, 1\}, & \text{if } i \in \gamma. \end{cases}$$

In other terms, $\prod_{\underline{n}} A_i \cong \{0, 1, 2, 3\}^\alpha \times \{1, 2, 3\}^\beta \times \{0, 1\}^\gamma$. Assume the factors are rearranged if necessary so that this correspondence is on-the-nose equality. By definition, θ_i is Θ for $i \in \alpha$ and 0_{A_i} for $i \notin \alpha$. Thus, the i -th entry of $\mathcal{A}/\boldsymbol{\theta}$ is

$$A_i/\theta_i = \begin{cases} \{\{0, 1\}, \{2\}, \{3\}\}, & \text{if } i \in \alpha, \\ \{\{1\}, \{2\}, \{3\}\}, & \text{if } i \in \beta, \\ \{\{0\}, \{1\}\}, & \text{if } i \in \gamma. \end{cases}$$

The i -th value $f(i)$ of the quotient solution must belong to A_i/θ_i . Let $\alpha = \alpha_0 \cup \alpha_1$ be such that $f(i) = \{0, 1\}$ when $i \in \alpha_0$ and $f(i) \in \{\{2\}, \{3\}\}$ when $i \in \alpha_1$. Then,

$$f \in \{\{0, 1\}\}^{\alpha_0} \times \{\{2\}, \{3\}\}^{\alpha_1} \times \{\{1\}, \{2\}, \{3\}\}^\beta \times \{\{0\}, \{1\}\}^\gamma. \quad (7.3)$$

Note that $f(i)$ is a singleton for all $i \notin \alpha_0$, and recall that $\bigcup\{a\} = a$. Define $g \in \prod_{\underline{n}} A_i$ as follows: $g(i) = 0$ if $i \in \alpha_0 \cup \gamma$, and $g(i) = \bigcup f(i)$, if $i \in \alpha_1 \cup \beta$. We will prove $g \in \text{Sol}(\mathcal{J}, \underline{n})$.

Fix an arbitrary constraint $(\sigma, R) \in \mathcal{C}$, let $S := \text{im } \sigma$ be the set of indices in the scope of R , and let $S^c := \underline{n} - S$ denote the complement of S with respect to \underline{n} . Since we assumed f solves \mathcal{J}/θ and satisfies (7.3) there must exist $\mathbf{r} \in \prod_{\underline{n}} A_i$ satisfying not only (σ, R) but also the following: $\mathbf{r}(i) \in \{0, 1\}$ for $i \in \alpha_0$ and $\mathbf{r}(i) = \bigcup f(i)$ for $i \notin \alpha_0$.

Now consider other elements satisfying the (arbitrary) constraint (σ, R) . By subdirectness of R , for each $\ell \in \alpha_0 \cup \gamma$ there exists $\mathbf{x}^\ell \in \prod_{\underline{n}} A_i$ satisfying (σ, R) and $\mathbf{x}^\ell(\ell) = 0$. Since \mathbf{r} and \mathbf{x}^ℓ satisfy (σ, R) —that is $\mathbf{r} \circ \sigma \in R$ and $\mathbf{x}^\ell \circ \sigma \in R$ —we also have $t(\mathbf{r}, \mathbf{x}^\ell) \circ \sigma \in R$. Consider the entries of $t(\mathbf{r}, \mathbf{x}^\ell)$, using the partition $\underline{n} = \alpha_0 \cup \alpha_1 \cup \beta \cup \gamma$ defined above. Since $\mathbf{x}^\ell(\ell) = 0$, we have $t(\mathbf{r}, \mathbf{x}^\ell)(\ell) = 0$. For $i \neq \ell$, observe the following:

- If $i \in \alpha_0 - \{\ell\}$, then $\mathbf{r}(i) \in \{0, 1\}$, so $t(\mathbf{r}, \mathbf{x}^\ell)(i) \in \{0, 1\}$.
- If $i \in \alpha_1$, then $\mathbf{r}(i) = \bigcup f(i) \in \{2, 3\}$, so $t(\mathbf{r}, \mathbf{x}^\ell)(i) = \bigcup f(i)$.
- If $i \in \beta$, then $\mathbf{r}(i) = \bigcup f(i) \in \{1, 2, 3\}$, so $t(\mathbf{r}, \mathbf{x}^\ell)(i) = \bigcup f(i)$.
- If $i \in \gamma - \{\ell\}$, then $\mathbf{r}(i) = \bigcup f(i) \in \{0, 1\}$, so $t(\mathbf{r}, \mathbf{x}^\ell)(i) \in \{0, 1\}$.

To summarize, for each $\ell \in \alpha_0 \cup \gamma$, there exists $t(\mathbf{r}, \mathbf{x}^\ell) \in \prod_{\underline{n}} A_i$ satisfying (σ, R) and having values $t(\mathbf{r}, \mathbf{x}^\ell)(i) \in \{0, 1\}$ for $i \in \alpha_0 \cup \gamma$ and

$$t(\mathbf{r}, \mathbf{x}^\ell)(i) = \begin{cases} 0, & i = \ell, \\ \bigcup f(i), & i \in \alpha_1 \cup \beta. \end{cases}$$

Finally, take the product of all members of $\{t(\mathbf{r}, \mathbf{x}^\ell) \mid \ell \in \alpha_0 \cup \gamma\}$ with respect to the binary operation of the original algebra $\langle A, \cdot \rangle$ to obtain $t(\mathbf{r}, \mathbf{x}^{\ell_1}) \cdot t(\mathbf{r}, \mathbf{x}^{\ell_2}) \cdots t(\mathbf{r}, \mathbf{x}^{\ell_L}) = g$, where $L = |\alpha_0 \cup \gamma|$. This proves that $g \in \text{Sol}((\sigma, R), \underline{n})$. Since the constraint (σ, R) was arbitrary, we have $g \in \text{Sol}(\mathcal{J}, \underline{n})$, as desired. \square

Example 7.4 and Proposition 7.6 do not fully illustrate the scope of the quotient-block strategy since, as we saw in the proof, every solution f to the quotient instance \mathcal{J}/θ induced a solution g on the block instance f/θ . In general, we may not be so lucky. Nonetheless, we demonstrated the quotient-block strategy on this example because the analysis carried out in Section 8 requires that the associated CSP is tractable (as asserted in Proposition 7.6).

7.4. The least block algorithm. This section describes an algorithm that was apparently discovered by “some subset of” Petar Marković and Ralph McKenzie. We first learned about this by reading Marković’s slides [sJM⁺11] from a 2011 talk in Krakow which gives a one-slide description of the algorithm. Since no proof (or even formal statement of the assumptions) from the originators of this result seems to be forthcoming, we include our own version and proof in this section.

Let \mathbf{A} be a finite idempotent algebra and let θ be a congruence on \mathbf{A} . Recall that by idempotence, every θ -class of A will be a subalgebra of \mathbf{A} . Suppose that, as an algebra, each θ -class comes from a variety, \mathcal{E} , possessing a $(k+1)$ -ary edge term, e . Define $s(x, y) = e(y, y, x, x, \dots, x)$. Then according to Definition 6.4, $\mathcal{E} \models s(x, y) \approx x$. By iterating s on its second variable, we obtain a binary term t such that

$$\mathbf{A} \models t(x, t(x, y)) \approx t(x, y) \tag{7.4}$$

$$\mathcal{E} \models t(x, y) \approx x. \tag{7.5}$$

Assume, finally, that the induced algebra $\langle A/\theta, t^{\mathbf{A}/\theta} \rangle$ is a linearly ordered meet semilattice. When this occurs, we call \mathbf{A} an *Edge-by-Chain* (EC) algebra, and we say, “ \mathbf{A} is EC.” Our objective in this subsection is to prove the following theorem.

Theorem 7.7. *Every finite, idempotent EC algebra is tractable.*

Proof. Recall, an algebra \mathbf{A} is *tractable* if the problem $\text{CSP}(\mathbf{S}(\mathbf{A}))$ is solvable in polynomial time. (See Subsection 6.2.)

Let \mathbf{A} be finite, idempotent and EC. We continue to use the congruence θ and term t developed above. Let \mathcal{J} be an instance of the problem $\text{CSP}(\mathbf{S}(\mathbf{A}))$. Thus \mathcal{J} is a triple $\langle \mathcal{V}, \mathcal{A}, \mathcal{C} \rangle$ in which $\mathcal{A} = (\mathbf{A}_0, \mathbf{A}_1, \dots, \mathbf{A}_{n-1})$ is a list of subalgebras of \mathbf{A} , see Definition 6.1. The congruence θ induces a congruence (which we will continue to call θ) on each \mathbf{A}_i , and the term t will induce a linearly ordered semilattice structure on A_i/θ . In other words, each \mathbf{A}_i is also an EC algebra.

Therefore, for each $i < n$, the structure of $\langle A_i/\theta, t \rangle$ will be of the form $C_{i,0} < C_{i,1} < \dots < C_{i,q_i}$ in which each $C_{i,j}$ is a θ -class of \mathbf{A}_i . We have the following useful relationship.

$$i < n, j \leq k, u \in C_{i,j}, v \in C_{i,k} \implies t(u, v) = u. \quad (7.6)$$

To see this, let $w = t(u, v)$. Then $w \in C_{i,j}$, since the quotient structure is a semilattice under t . Therefore, by (7.4), $t(u, w) = t(u, t(u, v)) = t(u, v) = w$. On the other hand, since $u, w \in C_{i,j} \in \mathcal{E}$ we have $t(u, w) = u$ by (7.5). Thus $u = w$.

Recall that for an index $k \leq n$, the k th-partial instance of \mathcal{J} is the instance, \mathcal{J}_k , obtained by restricting \mathcal{J} to the first k variables (Section 7.1.2). The set of solutions to \mathcal{J}_k is denoted $\text{Sol}(\mathcal{J}, k)$.

For every $k < n$ we recursively define the index j_k by

$$j_k = \text{least } j \leq q_k \text{ such that } \text{Sol}(\mathcal{J}, k+1) \cap \prod_{i=1}^k C_{i,j_i} \neq \emptyset.$$

If there is a k such that no such j_k exists, then j_k, j_{k+1}, \dots are undefined.

To complete the proof of Theorem 7.7, we need to prove two final claims.

Claim 7.8. \mathcal{J} has a solution if and only if j_{n-1} is defined.

Claim 7.9. j_i can be computed in polynomial time for all $i < n$.

One direction of Claim 7.8 is immediate: if j_{n-1} is defined then $\text{Sol}(\mathcal{J}) = \text{Sol}(\mathcal{J}, n)$ is nonempty. We address the converse. Assume that \mathcal{J} has solutions. We argue, by induction on k , that every j_k is defined. For the base step, choose any $g \in \text{Sol}(\mathcal{J})$. Since $g(0) \in A_0$, there is some ℓ such that $g(0) \in C_{0,\ell}$. Then j_0 is defined and is at most ℓ .

Now assume that j_0, j_1, \dots, j_{k-1} are defined, but j_k is undefined. Let $f \in \text{Sol}(\mathcal{J}, k) \cap \prod_{i=0}^{k-1} C_{i,j_i}$. Since j_k is undefined, f has no extension to a member of $\text{Sol}(\mathcal{J}, k+1)$. We shall derive a contradiction. Since \mathcal{J} has a solution, the restricted instance \mathcal{J}_{k+1} certainly has solutions. Choose any $g \in \text{Sol}(\mathcal{J}, k+1)$ with $g(k) \in C_{k,\ell}$ for the smallest possible ℓ . Let $\mathcal{C} = ((\sigma_0, R_0), (\sigma_1, R_1), \dots, (\sigma_{J-1}, R_{J-1}))$ be the list of constraints of \mathcal{J} . By assumption, for each $m < J$, f is a solution to $(\sigma_m, R_m)|_k$ (the restriction of R_m to its first k variables). Therefore, there is an extension, f_m , of f , to $k+1$ variables such that f_m is a solution to $(\sigma_m, R_m)|_{k+1}$.

Let $m < J$ and set $h_m = t(g, f_m)$. Since both g and f_m are solutions to $(\sigma_m, R_m)|_{k+1}$, and t is a term, h_m is also a solution to $(\sigma_m, R_m)|_{k+1}$. For every $i < k$ we have $f_m(i) = f(i) = f_0(i)$. Thus $h_m(i) = t(g(i), f_m(i)) = t(g(i), f_0(i)) = h_0(i)$ for $i < k$, and $h_m(k) = t(g(k), f_m(k)) = g(k)$. The latter relation holds by (7.6) and the choice of g . It follows

that every h_m coincides with h_0 . Since h_m satisfies R_m , we have $h_0 \in \text{Sol}(\mathcal{J}, \underline{k+1})$. This contradicts our assumption that f has no extension to a member of $\text{Sol}(\mathcal{J}, \underline{k+1})$.

Finally, we must verify Claim 7.9. Assume we have computed j_i , for $i < k$. To compute j_k we proceed as follows.

```

for  $j = 0$  to  $q_k$  do
   $\mathcal{B} \leftarrow (\mathbf{C}_{0,j_0}, \mathbf{C}_{1,j_1}, \dots, \mathbf{C}_{k-1,j_{k-1}}, \mathbf{C}_{k,j})$ 
  if  $\mathcal{J}_{\mathcal{B}}$  has a solution in  $\text{CSP}(\mathcal{B})$  then return  $j$ 
return FAILURE

```

In this algorithm, $\mathcal{J}_{\mathcal{B}}$ is the instance of $\text{CSP}(\mathcal{B})$ obtained by first restricting \mathcal{J} to its first $k+1$ variables, and then restricting each constraint relation to $\prod \mathcal{B}$. Since the members of \mathcal{B} all lie in the edge-term variety \mathcal{E} , it follows from Theorems 6.5 and 6.6 that $\text{CSP}(\mathcal{B})$ runs in polynomial-time. \square

8. CSPs OF COMMUTATIVE IDEMPOTENT BINARS

Taylor terms were defined in Section 2.2.2. It is not hard to see that a binary term is Taylor if and only if it is idempotent and commutative. This suggests, in light of the algebraic CSP-dichotomy conjecture, that we study commutative, idempotent binars (CIB's for short), that is, algebras with a single basic binary operation that is commutative and idempotent. If the conjecture is true, then every finite CIB should be tractable.

The associative CIBs are precisely the semilattices. Finite semilattices have long been known to be tractable, [JCG97]. The variety of semilattices is $\text{SD-}\wedge$ and is equationally complete. Every nontrivial semilattice must contain a subalgebra isomorphic to the unique two-element semilattice, $\mathbf{S}_2 = \langle \{0, 1\}, \wedge \rangle$, and conversely, \mathbf{S}_2 generates the variety of all semilattices.

As we shall see, the algebra \mathbf{S}_2 plays a central role in the structure of CIBs. In particular, the omission of \mathbf{S}_2 implies tractability.

Theorem 8.1. *If \mathbf{A} is a finite CIB then the following are equivalent:*

- (1) \mathbf{A} has an edge term.
- (2) $\mathcal{V}(\mathbf{A})$ is congruence modular.
- (3) $\mathbf{S}_2 \notin \text{HS}(\mathbf{A})$

The proof that (1) implies (2) appears in [BIM⁺10] and holds for general finite idempotent algebras. The contrapositive of (2) implies (3) is easy: if $\mathbf{S}_2 \in \text{HS}(\mathbf{A})$, then $\mathbf{S}_2^2 \in \mathcal{V}(\mathbf{A})$, and the congruence lattice of \mathbf{S}_2^2 is not modular. So it remains to show that (3) implies (1). For this we use the idea of a ‘‘cube-term blocker’’ ([MMM12]). A *cube term blocker* (CTB) for an algebra \mathbf{A} is a pair (D, S) of subuniverses with the following properties: $\emptyset < D < S \leq A$ and for every term $t(x_0, x_1, \dots, x_{n-1})$ of \mathbf{A} there is an index $i \in \underline{n}$ such that, for all $\mathbf{s} = (s_0, s_1, \dots, s_{n-1}) \in S^n$, if $s_i \in D$ then $t(\mathbf{s}) \in D$.

Theorem 8.2 ([MMM12, Thm 2.1]). *Let \mathbf{A} be a finite idempotent algebra. Then \mathbf{A} has an edge term iff it possesses no cube-term blockers.*

(The notions of cube term and edge term both originate in [BIM⁺10]. In that paper it is proved that a finite algebra has a cube term if and only if it has an edge term.)

Lemma 8.3. *A finite CIB \mathbf{A} has a CTB if and only if $\mathbf{S}_2 \in \text{HS}(\mathbf{A})$.*

Proof. Assume (D, S) is a CTB for \mathbf{A} . Then there exists $s \in S - D$. Consider $D^+ := D \cup \{s\}$. Evidently D^+ is a subuniverse of \mathbf{A} , and if \mathbf{D}^+ denotes the corresponding subalgebra, then $\theta_D := D^2 \cup \{(s, s)\}$ is a congruence of \mathbf{D}^+ . It's easy to see that $\mathbf{D}^+ / \theta_D \cong \mathbf{S}_2$, so $\mathbf{S}_2 \in \text{HS}(\mathbf{A})$.

Conversely, if $\mathbf{S}_2 \in \text{HS}(\mathbf{A})$, then there exists $\mathbf{B} \leq \mathbf{A}$ and a surjective homomorphism $h: \mathbf{B} \rightarrow \mathbf{S}_2$. Let $B_0 = h^{-1}(0)$. Then $\emptyset \neq B_0 < B \leq A$, and (B_0, B) is a CTB for \mathbf{A} . \square

Lemma 8.3, along with Theorem 8.2, completes the proof of Theorem 8.1 by showing (1) is false if and only if (3) is false. In fact, something stronger is true. Kearnes has shown that if \mathcal{V} is any variety of CIBs that omits \mathbf{S}_2 , then \mathcal{V} is congruence permutable [Kea16].

Before proving the next corollary, we need two more definitions. A variety is called *regular* if it is defined by regular identities. In contrast, a variety is called *strongly irregular* if it satisfies an identity $t(x, y) \approx x$ for some binary term t in which both x and y appear [BF15]. Every strongly irregular variety has an equational base consisting of a set of regular identities and a single strongly irregular identity. Note that most “interesting” varieties are strongly irregular—most Maltsev conditions involve a strongly irregular identity. For example, the Maltsev condition for congruence-permutability has a ternary term $q(x, y, z)$ satisfying $q(x, y, y) \approx x$, which is a strongly irregular identity. By contrast, the variety of semilattices is regular.

The proof of the next result will refer to *strongly irregular varieties* and *Plonka sums*. Our definitions of these concepts are the same as those found in [BF15], but we include them here for completeness. A variety is called *strongly irregular* if it satisfies an identity $t(x, y) \approx x$ for some binary term t in which both x and y appear. A similarity type of algebras is said to be *plural* if it contains no nullary operation symbols, and at least one non-unary operation symbol.

Let $\langle S, \vee \rangle$ be a semilattice, $\{\mathbf{A}_s \mid s \in S\}$ a collection of algebras of plural type $\rho: \mathcal{F} \rightarrow \mathbb{N}$, and $\{\varphi_{s,t}: \mathbf{A}_s \rightarrow \mathbf{A}_t \mid s \leq_{\vee} t\}$ a collection of homomorphisms satisfying $\varphi_{s,s} = 1_{A_s}$ and $\varphi_{t,u} \circ \varphi_{s,t} = \varphi_{s,u}$. The *Plonka sum* of the system $\langle \mathbf{A}_s : s \in S; \varphi_{s,t} : s \leq_{\vee} t \rangle$ is the algebra \mathbf{A} of type ρ with universe $A = \bigcup \{\mathbf{A}_s \mid s \in S\}$ and for $f \in \mathcal{F}$ a basic n -ary operation,

$$f^{\mathbf{A}}(x_1, x_2, \dots, x_n) = f^{\mathbf{A}_s}(\varphi_{s_1, s}(x_1), \varphi_{s_2, s}(x_2), \dots, \varphi_{s_n, s}(x_n))$$

in which $s = s_1 \vee s_2 \vee \dots \vee s_n$ and $x_i \in A_{s_i}$ for $1 \leq i \leq n$.

Corollary 8.4. *Let $\mathbf{A}_0, \mathbf{A}_1, \dots, \mathbf{A}_{n-1}$ be finite CIBs satisfying the equivalent conditions in Theorem 8.1. Then both $\mathbf{A}_0 \times \dots \times \mathbf{A}_{n-1}$ and $\mathbf{A}_0 \times \dots \times \mathbf{A}_{n-1} \times \mathbf{S}_2$ are tractable.*

Proof. By the theorem, $\mathcal{V}(\mathbf{A}_0)$ and $\mathcal{V}(\mathbf{A}_1)$ each have an edge term. By Theorem 1.2, $\mathbf{H}(\mathcal{V}(\mathbf{A}_0) \circ \mathcal{V}(\mathbf{A}_1))$ has an edge term as well. Since $\mathbf{A}_0 \times \mathbf{A}_1$ lies in this variety, it possesses that same edge term. Iterating this process, the algebra $\mathbf{B} = \mathbf{A}_0 \times \dots \times \mathbf{A}_{n-1}$ has an edge term, hence is tractable by Theorem 6.5.

Let $t(x_1, \dots, x_{k+1})$ be an edge term for \mathbf{B} . Then according to the first identity in Definition 6.4, $\mathcal{V}(\mathbf{B})$ is a strongly irregular variety. Consequently, the algebra $\mathbf{B} \times \mathbf{S}_2$ is a Plonka sum of two copies of \mathbf{B} . As a Plonka sum of tractable algebras, Theorem 4.1 of [BF15] implies that $\mathbf{B} \times \mathbf{S}_2$ is tractable as well. \square

It follows from the corollary and Theorem 6.6 that every finite member of $\mathcal{V}(\mathbf{A}_0 \times \mathbf{A}_1 \times \dots \times \mathbf{A}_{n-1} \times \mathbf{S}_2)$ is at least locally tractable.

As a counterpoint to Theorem 8.1, we mention the following. The proof requires some basic tame congruence theory, for which we refer the reader to [HM88, Thm 9.10].

Theorem 8.5. *Let \mathbf{A} be a finite CIB and suppose that $\text{HS}(\mathbf{A})$ contains no nontrivial abelian algebras. Then $\mathcal{V}(\mathbf{A})$ is SD- \wedge . Consequently, \mathbf{A} is a tractable algebra.*

Proof. Since every CIB has a Taylor term, and since $\text{HS}(\mathbf{A})$ has no abelian algebras, $\mathbf{S}(\mathbf{A})$ omits types $\{1, 2\}$. Then by [FV09, Cor 2.2], $\mathcal{V}(\mathbf{A})$ omits types $\{1, 2\}$. But then, by [HM88, Thm 9.10], $\mathcal{V}(\mathbf{A})$ is SD- \wedge . \square

In a congruence-permutable variety, an algebra is abelian if and only if it is polynomially equivalent to a faithful, unital module over a ring. See [Ber12, Thm 7.35] for a full discussion. For example, consider the ring \mathbb{Z}_3 as a module over itself. Define the binary polynomial $x \cdot y = 2x + 2y$ on this module. The table for this operation is given at the left of Figure 2. Conversely, we can retrieve the module operations from ‘ \cdot ’ by $x + y = 0 \cdot (x \cdot y)$ and $2x = 0 \cdot x$. Consequently, the commutative idempotent binar \mathbf{Sq}_3 is abelian.

In light of Theorem 6.7, every finite, abelian CIB will be of this form. We make the following observation.

Proposition 8.6. *A finite, abelian CIB has odd order.*

Proof. Let \mathbf{A} be a finite, abelian CIB. By Theorem 6.7 (since every CIB is Taylor), \mathbf{A} is polynomially equivalent to a faithful, unital module M over a ring R . The basic operation of \mathbf{A} must be a polynomial of M . Thus there are $r, s \in R$ and $b \in M$ such that $x \cdot y = rx + sy + b$. Let 0 denote the zero element of M . Then $0 = 0 \cdot 0 = r0 + s0 + b = b$. The condition $x \cdot y = y \cdot x$ implies $rx + sy = ry + sx$, so (by taking $y = 0$ and since M is faithful), $r = s$. Finally by idempotence, $2r = 1$.

Now, if A , hence M , has even cardinality, there is an element $a \in A$, $a \neq 0$, of additive order 2. Thus $a = a \cdot a = 2ra = 0$ which is a contradiction. \square

To demonstrate the utility of the techniques developed in this paper, we now show that every CIB of cardinality at most 4 is tractable. It is already known that the algebraic dichotomy conjecture holds for every idempotent algebra of cardinality at most 3 [Bul06, Sch78]. However, our arguments are relatively short and help us better understand CIBs and their distinctive properties.

For the remainder of this section, let \mathbf{A} be a CIB with universe n . We consider the various possibilities for n and \mathbf{A} (up to isomorphism), and in each case show that the associated CSP is tractable. In order to make it clear that our inventory is complete, a few of the arguments and computations are postponed until the end.

$n = 2$. There is a unique 2-element CIB, namely \mathbf{S}_2 . Idempotence determines the diagonal entries in the table, and by commutativity, the remaining two entries must be equal. The choices 0 and 1 for the off-diagonal yield a meet-semilattice and a join-semilattice respectively. As we remarked above, every finite semilattice is tractable, and, in fact, Sl , the variety of semilattices, is SD- \wedge .

$n = 3$, \mathbf{A} not simple. Let $\theta \in \text{Con } \mathbf{A}$ be nontrivial (i.e., $0_A < \theta < 1_A$). By idempotence, every congruence class is a subalgebra. Evidently, \mathbf{A}/θ has cardinality 2, one θ class has 2 elements, and the other has 1. From the uniqueness of the 2-element algebra, \mathbf{A}/θ and each of the θ -classes are semilattices. Consequently, $\mathbf{A} \in Sl \circ Sl$. By Theorem 1.2, $Sl \circ Sl$ is SD- \wedge , so by Theorem 6.3, \mathbf{A} is tractable.

$n = 3$, \mathbf{A} simple. If \mathbf{A} has no proper nontrivial subalgebras, then in \mathbf{A} , $x \neq y \implies x \cdot y \notin \{x, y\}$. It follows that \mathbf{A} must be the 3-element Steiner quasigroup, \mathbf{Sq}_3 , of Figure 2. By Corollary 8.4, \mathbf{A} is tractable.

·	0	1	2
0	0	2	1
1	2	1	0
2	1	0	2

\mathbf{Sq}_3

·	0	1	2
0	0	0	1
1	0	1	2
2	1	2	2

\mathbf{T}_1

·	0	1	2
0	0	0	2
1	0	1	1
2	2	1	2

\mathbf{T}_2

Figure 2: The simple CIBs of cardinality 3

Finally, if \mathbf{A} has a proper nontrivial subalgebra, that subalgebra must be isomorphic to \mathbf{S}_2 . Thus no nontrivial subalgebra of \mathbf{A} is abelian. By Theorem 8.5, \mathbf{A} generates an SD- \wedge variety, so is tractable by Theorem 6.3. For future reference, we remark that there are two non-isomorphic algebras meeting the description in this paragraph, \mathbf{T}_1 and \mathbf{T}_2 . Their tables are given in Figure 2.

$n = 4$, \mathbf{A} not simple. Let θ be a maximal congruence of \mathbf{A} . Then \mathbf{A}/θ is simple, so \mathbf{A}/θ is isomorphic to one of \mathbf{T}_1 , \mathbf{T}_2 , \mathbf{Sq}_3 , or \mathbf{S}_2 . If $\mathbf{A}/\theta \cong \mathbf{T}_i$, then the θ -classes must have sizes 1, 1, and 2. Consequently, each θ -class is a semilattice, so $\mathbf{A} \in Sl \circ \mathcal{V}(\mathbf{T}_i)$ which is SD- \wedge by the computations above and Theorem 1.2. Thus \mathbf{A} is tractable.

The next case to consider is $\mathbf{A}/\theta \cong \mathbf{Sq}_3$. Without loss of generality, assume that θ partitions the universe as $|01|2|3|$, and further that $0 \cdot 1 = 0$. From these data we deduce that the operation table for \mathbf{A} must be

·	0	1	2	3
0	0	0	3	2
1	0	1	3	2
2	3	3	2	a
3	2	2	a	3

with $a \in \{0, 1\}$.

If $a = 0$, then the only way to obtain 1 as a product is $1 \cdot 1$. In that case there is a homomorphism onto \mathbf{S}_2 with kernel $\psi = |023|1|$. We have $\theta \cap \psi = 0_A$, so \mathbf{A} is a subdirect product of $\mathbf{Sq}_3 \times \mathbf{S}_2$. Therefore, by Corollary 8.4, \mathbf{A} is tractable. More problematic is the case $a = 1$, but we already established tractability of that algebra in Example 7.4.

Finally, suppose that $\mathbf{A}/\theta \cong \mathbf{S}_2$. The possible sizes of the θ -classes are 3, 1 or 2, 2. If neither class is isomorphic to \mathbf{Sq}_3 then the classes both lie in an SD- \wedge variety, so by Theorem 1.2, \mathbf{A} too lies in an SD- \wedge variety and is thus tractable. If one of the classes is isomorphic to \mathbf{Sq}_3 (there are 7 such algebras), then \mathbf{A} is an EC algebra (the term $t(x, y) = y \cdot (x \cdot y)$ is Maltsev, hence is an edge term) and we can apply Theorem 7.7 to establish that \mathbf{A} is tractable.

$n = 4$, \mathbf{A} simple. By Theorem 8.5, if $\mathbf{HS}(\mathbf{A})$ contains no nontrivial abelian algebra, then \mathbf{A} is tractable. Thus, we may as well assume that $\mathbf{HS}(\mathbf{A})$ contains an abelian algebra. By Proposition 8.6, \mathbf{A} itself is nonabelian. Consequently (since \mathbf{A} is simple) \mathbf{A} must have a subalgebra isomorphic to \mathbf{Sq}_3 . Without loss of generality, let us assume that $\{1, 2, 3\}$ forms this subalgebra. Similarly, by Corollary 8.4, we can assume that $\mathbf{HS}(\mathbf{A})$ contains a copy of \mathbf{S}_2 . Examining the 2- and 3-element CIBs, it is easy to see that if $\mathbf{S}_2 \in \mathbf{HS}(\mathbf{A})$, then already $\mathbf{S}_2 \in \mathbf{S}(\mathbf{A})$. By the symmetry of \mathbf{Sq}_3 , and the fact that it contains no semilattice, we can assume that $\{0, 1\}$ forms the semilattice. Thus, the table for \mathbf{A} must have one of

	u_2	u_3	Proper Nontrivial Subalgebras
\mathbf{A}_0	0	1	$\{0, 1\}, \{0, 2\}, \{1, 2, 3\}$
\mathbf{A}_1	1	1	$\{0, 1\}, \{1, 2, 3\}$
\mathbf{A}_2	1	2	$\{0, 1\}, \{1, 2, 3\}$
\mathbf{A}_3	0	3	$\{0, 1\}, \{0, 2\}, \{0, 3\}, \{1, 2, 3\}$
\mathbf{A}_4	1	3	$\{0, 1\}, \{0, 3\}, \{1, 2, 3\}$
\mathbf{A}_5	2	2	$\{0, 1\}, \{0, 2\}, \{0, 3\}, \{1, 2, 3\}$
\mathbf{A}_6	2	3	$\{0, 1\}, \{0, 2\}, \{0, 3\}, \{1, 2, 3\}$

Figure 3: The 7 simple CIBs of size 4.

the following two forms.

\cdot	0	1	2	3	\cdot	0	1	2	3
0	0	0	u_2	u_3	0	0	1	v_2	v_3
1	0	1	3	2	1	1	1	3	2
2	u_2	3	2	1	2	v_2	3	2	1
3	u_3	2	1	3	3	v_3	2	1	3

By checking the 32 possible tables, either via the universal algebra calculator [FKV11] or directly using Freese's algorithm [Fre08], one determines that there are 7 pairwise non-isomorphic algebras of one of these two forms. Interestingly, every simple algebra of the second form is isomorphic to one of the first form. Thus we will use the first form for all the candidates. These 7 algebras are indicated in Figure 3.

Our intent is to apply the rectangularity theorem to establish the tractability of these 7 algebras. For this, we need to determine the minimal absorbing subuniverses of each subalgebra of \mathbf{A} . First observe that if $\{a, b\}$ forms a copy of \mathbf{S}_2 with $a \cdot b = a$, then $\{a\} \triangleleft_t \{a, b\}$, with $t(x, y) = x \cdot y$, while $\{b\}$ is not absorbing since a is a sink (Lemma 4.10). Second, by Lemma 4.13, \mathbf{Sq}_3 is absorption-free, which is to say, it is its own minimal absorbing subalgebra.

Now let B be a proper minimal absorbing subalgebras of \mathbf{A}_i , for $i < 7$. Then either $B \cap \{1, 2, 3\}$ is absorbing in $\{1, 2, 3\}$ or it is empty (Lemma 4.8). As $\{1, 2, 3\} \cong \mathbf{Sq}_3$ is absorption free and B is a proper subalgebra of \mathbf{A} , we must have $B = \{1, 2, 3\}$ or $B = \{0\}$. However the first of these is impossible because $\{1, 2, 3\} \cap \{0, 1\} = \{1\}$, which is not absorbing in $\{0, 1\}$ since 0 is a sink. Thus the only possible proper absorbing subalgebra of \mathbf{A} is $\{0\}$.

For $i \leq 2$, $\{0\}$ is indeed absorbing in \mathbf{A}_i . We believe that the shortest binary absorbing term in each case is as follows.

$$\begin{aligned}
\mathbf{A}_0 & (x(xy))(y(xy)) \\
\mathbf{A}_1 & (x(x(xy)))(y(y(xy))) \\
\mathbf{A}_2 & ((x(x(xy)))(y(xy)(x(x(xy)))))(y(x(xy))(x(x(x(xy))))).
\end{aligned} \tag{8.1}$$

However, observe that if $i > 2$ then according to Figure 3 either $u_2 = 2$ or $u_3 = 3$. If $u_2 = 2$ then 2 is a sink for the subuniverse $\{0, 2\}$. In that case $\{0\} = B \cap \{0, 2\}$ is not absorbing in $\{0, 2\}$ contradicting the fact that B is absorbing in \mathbf{A} . A similar argument works for $u_3 = 3$ and $\{0, 3\}$. Thus \mathbf{A}_i is absorption free for $i > 2$.

We summarize these computations in the following table.

Algebra	Minimal Absorbing Subalgebra
$\mathbf{S}_2 = \{0, 1\}$	$\{0\}$
$\{1, 2, 3\}$	$\{1, 2, 3\}$
$\mathbf{A}_0, \mathbf{A}_1, \mathbf{A}_2$	$\{0\}$
$\mathbf{A}_3, \dots, \mathbf{A}_6$	A

The crucial point here is that every member of $\mathbf{S}(\mathbf{A}_i)$ has a unique minimal absorbing subalgebra.

We make one more observation. Let $i < 7$ and let h be an automorphism of \mathbf{A}_i . Then $h(0) = 0$. This holds because $\{1, 2, 3\}$ must be preserved by h , as it is the unique 3-element subalgebra. Of course since \mathbf{S}_2 has no nontrivial automorphisms, $h(0) = 0$ also holds in that algebra.

We now proceed to argue that each algebra from Figure 3 is tractable. Following Definition 6.1, let $\mathcal{J} = \langle \mathcal{V}, \mathcal{A}, \mathcal{S}, \mathcal{R} \rangle$ be an instance of $\text{CSP}(\mathbf{A}_i)$, for some $i < 7$. We shall simply write \mathbf{A} instead of \mathbf{A}_i . We have $\mathcal{A} = (\mathbf{C}_0, \mathbf{C}_1, \dots, \mathbf{C}_{n-1})$, in which each \mathbf{C}_i is a subalgebra of \mathbf{A} , and $\mathcal{R} = (R_0, \dots, R_{J-1})$ is a list of subdirect products of the various \mathbf{C}_k 's. We must show that there is an algorithm that determines, in polynomial time, whether this instance has a solution. For this we will use Corollary 5.8.

For every $k < n$ let \mathbf{B}_k be the minimal absorbing subuniverse of \mathbf{C}_k . As mentioned, \mathbf{B}_k is unique. Let $\alpha = \{k < n : \mathbf{C}_k \text{ is abelian}\}$ and set $\alpha' = n - \alpha$. By reordering the factors, we may assume that $\alpha = \{0, 1, \dots, p-1\}$ for some $p < n$.

Fix $j < J$. There is an equivalence relation on α' defined by $k \sim_j \ell$ iff $\rho_k^j = \rho_\ell^j$. Call the least index in each equivalence class the *leader*. The remaining elements are the *followers*. When $k \sim_j \ell$, subdirectness of R_j implies that $\mathbf{C}_k \cong \mathbf{R}_j / \rho_k^j = \mathbf{R}_j / \rho_\ell^j \cong \mathbf{C}_\ell$. Thus, there is an isomorphism $h_{k,\ell}^j : \mathbf{C}_k \rightarrow \mathbf{C}_\ell$ such that $\text{Proj}_{k,\ell}(R_j) = \{(x, h_{k,\ell}^j(x)) : x \in \mathbf{C}_k\}$. These maps are consistent in the sense that if $k \sim \ell \sim m$ then $h_{k,m} = h_{\ell,m} \circ h_{k,\ell}$. Observe also that $h_{k,\ell}(B_k) = B_\ell$, since the minimal absorbing subalgebra is unique in each factor. Moreover, from our comments about automorphisms, and since $k, \ell \in \alpha'$, we have $h_{k,\ell}(0) = 0$. Finally, for $k \sim_j \ell$, define $H_{k,\ell}^j = \text{Proj}_{k,\ell}(R_j)$, so $(0, 0) \in H_{k,\ell}^j$ whenever $k \sim_j \ell$.

For each $j < J$, let the relation \bar{R}_j be obtained from R_j as follows. First, for each follower ℓ , remove the ℓ -th variable from the scope of R_j . Then, define $\bar{R}_j = R_j \times \prod_{\ell \notin \text{im } \sigma_j} C_\ell$. Let $\bar{\mathcal{R}} = (\bar{R}_j : j < J)$ and $\bar{\mathcal{H}} = (H_{k,\ell}^j : j < J, k \sim_j \ell)$. Lastly, let $\bar{\mathcal{J}}$ denote the modification of \mathcal{J} with $\bar{\mathcal{R}} \cup \bar{\mathcal{H}}$ replacing \mathcal{R} and every scope equal to \mathcal{V} . Then the solution set of $\bar{\mathcal{J}}$ is identical to the solution set of the original instance \mathcal{J} . (To be more precise, the relations in $\bar{\mathcal{R}} \cup \bar{\mathcal{H}}$ are pp-definable (see[BKW17, Sec. 3.1]) from \mathcal{R} and conversely.)

Consider the partial instance $\bar{\mathcal{J}}_p$. This is an instance of the problem $\text{CSP}(\mathbf{Sq}_3)$, which is tractable. If the partial instance has no solution, then neither does \mathcal{J} , and we halt. So suppose that $\bar{\mathcal{J}}_p$ has a solution, \mathbf{x} . (We do not actually need to know \mathbf{x} , just that it exists.) Since \mathbf{Sq}_3 is absorption-free, $\mathbf{x} \in \prod_\alpha C_k = \prod_\alpha B_k$.

We now wish to apply Corollary 5.8 to show that

$$\{\mathbf{x}\} \times \prod_{k \in \alpha'} B_k \subseteq \bigcap \bar{\mathcal{R}}. \quad (8.2)$$

The first 3 conditions follow from the definition of α and the construction of the \bar{R}_j 's. The fifth condition holds because of the solution \mathbf{x} to the partial instance. We must address the fourth requirement.

Fix $j < J$. We shall show that \bar{R}_j intersects $\prod B_k$. We have two cases to consider. First, suppose that $i < 3$. In that case the minimal absorbing subalgebra of \mathbf{A} is $\{0\}$, with corresponding absorbing term given in (8.1). Think of that term as a nonassociative product. For every $\ell \in \alpha$, $B_\ell = \{1, 2, 3\} = C_\ell$, while for $\ell \in \alpha'$, $B_\ell = \{0\}$. Since \bar{R}_j is subdirect, for every $\ell \in \alpha'$ there is a tuple $\mathbf{r}^\ell \in R_j$ with $r_\ell^\ell = 0$. Let \mathbf{r} be the ‘‘product’’ (under the absorbing term) of all the \mathbf{r}^ℓ 's. (For definiteness, associate to the right.) Then, since $\{0\}$ is absorbing, $\mathbf{r}_\ell = 0 \in B_\ell$ for every $\ell \in \alpha'$, and for $\ell \in \alpha$, $\mathbf{r}_\ell \in C_\ell = B_\ell$. Thus the condition is satisfied.

The argument when $i \geq 3$ is not very different. In this case, let $\beta = \{\ell < n : \mathbf{C}_\ell = \mathbf{S}_2\}$. Note that if $\ell \notin \beta$ then $B_\ell = C_\ell$. As before, for $\ell \in \beta$ let \mathbf{r}^ℓ be a tuple with $\mathbf{r}_\ell^\ell = 0$, and let \mathbf{r} be the product. (This time simply use the basic binary operation to compute the product. It is, after all, an absorbing term for \mathbf{S}_2 .) Then for $\ell \in \beta$, $r_\ell = 0 \in B_\ell$, while, for $\ell \notin \beta$, $r_\ell \in C_\ell = B_\ell$.

Thus inclusion (8.2) holds. Let $\mathbf{0}$ denote a tuple of 0's indexed by α' , and set \mathbf{y} to be the concatenation of \mathbf{x} with $\mathbf{0}$. Then $\mathbf{y} \in \bigcap \bar{\mathcal{R}}$ by (8.2). On the other hand, for any $j < J$ and $k \sim_j \ell$, we have $(y_k, y_\ell) = (0, 0) \in H_{k,\ell}^j$. Thus \mathbf{y} is a solution to $\bar{\mathcal{J}}$.

9. ACKNOWLEDGEMENTS

The authors wish to thank the anonymous referee for identifying errors, suggesting corrections, and providing many other valuable suggestions for improvements. The authors also gratefully acknowledge the support of National Science Foundation grant no. 1500218.

APPENDIX A. MISCELLANEOUS PROOFS

A.1. Proof of Lemma 4.3. To simplify notation, let $ABB \cdots B := A \times B \times B \times \cdots \times B$ (the usual Cartesian product) and let $t = f \star g$. We handle the $j = 1$ case of the definition of absorption. (The general case is no harder, but the notation becomes tedious.) That is, we prove

$$t[ABB \cdots B] = f[g[ABB \cdots B] \times g[BB \cdots B] \times \cdots \times g[BB \cdots B]] \subseteq B.$$

If f is the absorbing term, then since B is a subalgebra, we have $g[BB \cdots B] \subseteq B$. Hence, $t[ABB \cdots B] \subseteq f[ABB \cdots B] \subseteq B$. On the other hand, if g is the absorbing term, then $t[ABB \cdots B] \subseteq f[BB \cdots B] \subseteq B$. \square

A.2. Proof of Corollary 4.6. Since $\mathbf{B}_0 \triangleleft_f \mathbf{A}_0$, it follows that $\mathbf{B}_0 \times \mathbf{A}_1 \triangleleft_f \mathbf{A}_0 \times \mathbf{A}_1$. Similarly $\mathbf{A}_0 \times \mathbf{B}_1 \triangleleft_g \mathbf{A}_0 \times \mathbf{A}_1$. Hence, $\mathbf{B}_0 \times \mathbf{B}_1 = (\mathbf{B}_0 \times \mathbf{A}_1) \cap (\mathbf{A}_0 \times \mathbf{B}_1)$ is absorbing in $\mathbf{A}_0 \times \mathbf{A}_1$ with respect to $f \star g$, by Lemma 4.5. \square

A.3. Proof of Lemma 4.7. In case $n = 2$, the fact that $\mathbf{B} \triangleleft_s \mathbf{A}$ follows directly from Corollary 4.6. We first extend this result to *minimal* absorbing subalgebras, still for $n = 2$, and then an easy induction argument will complete the proof for arbitrary finite n .

Assume $\mathbf{B}_0 \triangleleft_{t_0} \mathbf{A}_0$ and $\mathbf{B}_1 \triangleleft_{t_1} \mathbf{A}_1$. Then $\mathbf{B} \triangleleft_s \mathbf{A}$ (Cor. 4.6), where $\mathbf{B} := \mathbf{B}_0 \times \mathbf{B}_1$, $\mathbf{A} := \mathbf{A}_0 \times \mathbf{A}_1$, and $s = t_0 \star t_1$. To show $\mathbf{B}_0 \times \mathbf{B}_1$ is minimal absorbing, let \mathbf{S} be a proper subalgebra of $\mathbf{B}_0 \times \mathbf{B}_1$. By transitivity of absorption, it suffices to prove that \mathbf{S} is not absorbing in $\mathbf{B}_0 \times \mathbf{B}_1$.

Let t be a term of arity q . For $b \in B_1$, the set $S^{-1}b := \{b_0 \in B_0 \mid (b_0, b) \in S\}$ is a subuniverse of \mathbf{B}_0 . Since \mathbf{S} is a proper subalgebra, there exists $b^* \in B_1$ such that $S^{-1}b^*$ is not all of B_0 . Therefore, $S^{-1}b^*$ is a proper subuniverse of \mathbf{B}_0 , so $S^{-1}b^*$ is not absorbing in \mathbf{A}_0 (by minimality of \mathbf{B}_0). Consequently,

$$\begin{aligned} \exists x_i \in S^{-1}b^*, \quad \exists b \in B_0, \quad \exists j < q, \quad \exists b' \notin S^{-1}b^* \quad \text{such that} \\ t^{\mathbf{A}_0}(x_0, x_1, \dots, x_{j-1}, b, x_{j+1}, \dots, x_{q-1}) = b'. \end{aligned}$$

Thus, $t^{\mathbf{A}_0 \times \mathbf{A}_1}((x_0, b^*), \dots, (x_{j-1}, b^*), (b, b^*), (x_{j+1}, b^*), \dots, (x_{q-1}, b^*)) = (b', b^*) \notin S$ (since $b' \notin S^{-1}b^*$). Finally, because $(x_i, b^*) \in S$ for all i , and since t was an arbitrary term, \mathbf{S} is not absorbing in $\mathbf{B}_0 \times \mathbf{B}_1$.

Now fix $n > 2$ and assume the result holds when there are at most $n - 1$ factors. Let $\mathbf{B}' := \mathbf{B}_0 \times \dots \times \mathbf{B}_{n-2}$ and $\mathbf{A}' := \mathbf{A}_0 \times \dots \times \mathbf{A}_{n-2}$. By the induction hypothesis, $\mathbf{B}' \triangleleft_{s'} \mathbf{A}'$, and since $\mathbf{B}_{n-1} \triangleleft_{t_{n-1}} \mathbf{A}_{n-1}$ we have (by the $n = 2$ case) $\mathbf{B}' \times \mathbf{B}_{n-1} \triangleleft_s \mathbf{A}' \times \mathbf{A}_{n-1}$, where $s = s' \star t_{n-1}$. \square

A.4. Proof of Corollary 4.9. Let $\mathbf{A} = \prod_i \mathbf{A}_i$ and $\mathbf{B} = \prod_i \mathbf{B}_i$. By Lemma 4.7, $\mathbf{B} \triangleleft_t \mathbf{A}$, so the result follows from Lemma 4.8 if we put $C = R$. \square

This is proved Section A.5.

A.5. Proof of Lemma 4.11. If $R' \neq \emptyset$, then for $1 \leq i \leq n$ the projection $S_i := \text{Proj}_i R'$ is also nonempty. We want to show $S_i = B_i$. By minimality of $\mathbf{B}_i \triangleleft \mathbf{A}_i$ and by transitivity of absorption, it suffices to prove $\mathbf{S}_i \triangleleft \mathbf{B}_i$. Assume $\mathbf{B}_i \triangleleft \mathbf{A}_i$ with respect to t , say, $k = \text{ar}(t)$. Fix $s_1, \dots, s_k \in S_i$, $b \in B_1$, and $j \leq k$. Then $\tilde{b}_i := t^{\mathbf{A}_i}(s_1, \dots, s_{j-1}, b, s_{j+1}, \dots, s_k) \in B_i$, and we must show $\tilde{b}_i \in S_i$.

Since \mathbf{R} is subdirect, there exist $a_i \in A_i$ (for all $i \neq j$) such that

$$\mathbf{a}^* := (a_1, \dots, a_{j-1}, b, a_{j+1}, \dots, a_n) \in R.$$

Also, for each $1 \leq j \leq n$ there exist $s_1(j), \dots, s_k(j)$ in S_j such that for all $1 \leq \ell \leq k$ we have

$$\mathbf{s}_\ell := (s_\ell(1), \dots, s_\ell(i-1), s_\ell, s_\ell(i+1), \dots, s_\ell(n)) \in R'.$$

Since all these n -tuples belong to R , so does $t^{\mathbf{A}_1 \times \dots \times \mathbf{A}_n}(\mathbf{s}_1, \dots, \mathbf{s}_{j-1}, \mathbf{a}^*, \mathbf{s}_{j+1}, \dots, \mathbf{s}_k)$, and the latter reduces to $(\tilde{b}_1, \dots, \tilde{b}_n)$. Since $\mathbf{B}_i \triangleleft \mathbf{A}_i$, we have $(\tilde{b}_1, \dots, \tilde{b}_n) \in (B_1 \times \dots \times B_n)$. Thus, $(\tilde{b}_1, \dots, \tilde{b}_n) \in R \cap B_1 \times \dots \times B_n$, so $\tilde{b}_i \in S_i$, as desired. Of course, the same argument works for all $1 \leq i \leq n$. \square

REFERENCES

- [Bar15] Libor Barto. A rectangularity theorem for simple Taylor algebras. Open Problems in Universal Algebra, a Shanks workshop at Vanderbilt University, <http://www.math.vanderbilt.edu/~moorm10/shanks/>, May 2015. URL: <http://www.math.vanderbilt.edu/~moorm10/shanks/>.
- [Ber12] Clifford Bergman. *Universal algebra*, volume 301 of *Pure and Applied Mathematics (Boca Raton)*. CRC Press, Boca Raton, FL, 2012. Fundamentals and selected topics.
- [BF15] Clifford Bergman and David Failing. Commutative idempotent groupoids and the constraint satisfaction problem. *Algebra Universalis*, 73(3-4):391–417, 2015. URL: <http://dx.doi.org/10.1007/s00012-015-0323-6>, doi:10.1007/s00012-015-0323-6.
- [BIM⁺10] Joel Berman, Paweł Idziak, Petar Marković, Ralph McKenzie, Matthew Valeriote, and Ross Willard. Varieties with few subalgebras of powers. *Trans. Amer. Math. Soc.*, 362(3):1445–1473, 2010. URL: <http://dx.doi.org/10.1090/S0002-9947-09-04874-0>, doi:10.1090/S0002-9947-09-04874-0.
- [BJK05] Andrei Bulatov, Peter Jeavons, and Andrei Krokhin. Classifying the complexity of constraints using finite algebras. *SIAM J. Comput.*, 34(3):720–742, 2005. URL: <http://dx.doi.org/10.1137/S0097539700376676>, doi:10.1137/S0097539700376676.
- [BK09] Libor Barto and Marcin Kozik. Constraint satisfaction problems of bounded width. In *2009 50th Annual IEEE Symposium on Foundations of Computer Science (FOCS 2009)*, pages 595–603. IEEE Computer Soc., Los Alamitos, CA, 2009. URL: <http://dx.doi.org/10.1109/FOCS.2009.32>, doi:10.1109/FOCS.2009.32.
- [BK10] Libor Barto and Marcin Kozik. New conditions for Taylor varieties and CSP. In *25th Annual IEEE Symposium on Logic in Computer Science LICS 2010*, pages 100–109. IEEE Computer Soc., Los Alamitos, CA, 2010.
- [BK12] Libor Barto and Marcin Kozik. Absorbing subalgebras, cyclic terms, and the constraint satisfaction problem. *Log. Methods Comput. Sci.*, 8(1):1:07, 27, Feb 2012. URL: <https://lmcs.episciences.org/673>, doi:10.2168/LMCS-8(1:7)2012.
- [BK17] Libor Barto and Marcin Kozik. Absorption in Universal Algebra and CSP. In Andrei Krokhin and Stanislav Zivny, editors, *The Constraint Satisfaction Problem: Complexity and Approximability*, volume 7 of *Dagstuhl Follow-Ups*, pages 45–77. Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik, Dagstuhl, Germany, 2017. URL: <http://drops.dagstuhl.de/opus/volltexte/2017/6960>, doi:10.4230/DFU.Vol7.15301.45.
- [BKS15] Libor Barto, Marcin Kozik, and David Stanovský. Maltsev conditions, lack of absorption, and solvability. *Algebra Universalis*, 74(1-2):185–206, 2015. URL: <http://dx.doi.org/10.1007/s00012-015-0338-z>, doi:10.1007/s00012-015-0338-z.
- [BKW17] Libor Barto, Andrei Krokhin, and Ross Willard. Polymorphisms, and how to use them. In *The constraint satisfaction problem: complexity and approximability*, volume 7 of *Dagstuhl Follow-Ups*, pages 1–44. Schloss Dagstuhl. Leibniz-Zent. Inform., Wadern, 2017.
- [Bul06] Andrei A. Bulatov. A dichotomy theorem for constraint satisfaction problems on a 3-element set. *J. ACM*, 53(1):66–120, 2006. URL: <http://dx.doi.org/10.1145/1120582.1120584>, doi:10.1145/1120582.1120584.
- [Bul17] Andrei Bulatov. A dichotomy theorem for nonuniform CSPs. <https://arxiv.org/abs/1703.03021>, 2017.
- [FKV11] Ralph Freese, Emil Kiss, and Matthew Valeriote. Universal Algebra Calculator, 2011. Available at: <http://uacalc.org>.
- [FM17] Ralph Freese and Ralph McKenzie. Maltsev families of varieties closed under join or Maltsev product. *Algebra Universalis*, 77(1):29–50, 2017. URL: <https://doi.org/10.1007/s00012-016-0420-1>, doi:10.1007/s00012-016-0420-1.
- [Fre08] R. Freese. Computing congruences efficiently. *Alg. Univ.*, 59:337–343, 2008.
- [FV99] Tomás Feder and Moshe Y. Vardi. The computational structure of monotone monadic SNP and constraint satisfaction: a study through Datalog and group theory. *SIAM J. Comput.*, 28(1):57–104 (electronic), 1999. URL: <http://dx.doi.org/10.1137/S0097539794266766>, doi:10.1137/S0097539794266766.

- [FV09] Ralph Freese and Matthew A. Valeriote. On the complexity of some Maltsev conditions. *Internat. J. Algebra Comput.*, 19(1):41–77, 2009. URL: <http://dx.doi.org/10.1142/S0218196709004956>, doi:10.1142/S0218196709004956.
- [HM88] David Hobby and Ralph McKenzie. *The structure of finite algebras*, volume 76 of *Contemporary Mathematics*. American Mathematical Society, Providence, RI, 1988.
- [IMM⁺10] Paweł Idziak, Petar Marković, Ralph McKenzie, Matthew Valeriote, and Ross Willard. Tractability and learnability arising from algebras with few subpowers. *SIAM J. Comput.*, 39(7):3023–3037, 2010. URL: <http://dx.doi.org/10.1137/090775646>, doi:10.1137/090775646.
- [JCG97] Peter Jeavons, David Cohen, and Marc Gyssens. Closure properties of constraints. *J. ACM*, 44(4):527–548, 1997. URL: <http://dx.doi.org/10.1145/263867.263489>, doi:10.1145/263867.263489.
- [Kea16] Keith Kearnes. Personal communication, 2016.
- [KK13] Keith A. Kearnes and Emil W. Kiss. The shape of congruence lattices. *Mem. Amer. Math. Soc.*, 222(1046):viii+169, 2013. URL: <http://dx.doi.org/10.1090/S0065-9266-2012-00667-8>, doi:10.1090/S0065-9266-2012-00667-8.
- [MMM12] Petar Marković, Miklós Maróti, and Ralph McKenzie. Finitely related clones and algebras with cube terms. *Order*, 29(2):345–359, 2012. URL: <http://dx.doi.org/10.1007/s11083-011-9232-2>, doi:10.1007/s11083-011-9232-2.
- [Sch78] Thomas J. Schaefer. The complexity of satisfiability problems. In *Conference Record of the Tenth Annual ACM Symposium on Theory of Computing (San Diego, Calif., 1978)*, pages 216–226. ACM, New York, 1978.
- [sJM⁺11] B. Bašić, J. Jovanović, P. Marković, et al. CSP on small templates. <http://2oal.tcs.uj.edu.pl>, June 2011.
- [Tay77] Walter Taylor. Varieties obeying homotopy laws. *Canad. J. Math.*, 29(3):498–527, 1977.
- [Zhu17] Dmitriy Zhuk. A proof of CSP dichotomy conjecture. <https://arxiv.org/abs/1704.01914>, 2017.